Review: Network Monitoring Systems p. 67

Windows IT Pro

A PENTON PUBLICATION

NOVEMBER 2010 | WINDOWSITPRO.COM | WE'RE

Manage AD in Bulk with PowerShell p. 20

Exchange 2010 MRM: Modify and Reduce Help Desk Calls p. 47

Save Time with Windows Deployment Services p. 25

Master Scripting Shortcuts that Contain Unicode p. 51

Virtualization Features in Windows Server 2008 R2 Service Pack 1 p. 33

Configure Cross-Platform Support for System Center Ops Manager 2007 R2 p. 54

Inside the Ops Manager Management Pack p. 38

SharePoint Records Management: Design and Build a Compliant Platform p. 59

Extend the Active Directory Schema p. 43

Redefining X.

When an organization needs more computing power for today's memory-intensive workloads, the conventional wisdom is to buy more servers. This can lead to massive inefficiency and server sprawl, with the majority of servers today running at only 10% utilization.[1] As the computational demands of a smarter planet continue to explode, this sort of inefficiency has become a problem—a problem IBM engineers have now solved. The 5th generation of Enterprise X-Architecture® from IBM featuring the Intel® Xeon® Processor 7500 Series lets you add memory independently of the processor. As a result, IBM eX5 systems can leverage 6x more memory than current x86 servers, reduce storage costs by up to 97% and cut licensing fees by 50%.[2]

A smarter business needs smarter software, systems and services.
Let's build a smarter planet. ibm.com/systems/ex5

IBM®

intel Xeon inside™

Powerful.
Intelligent.

1. McKinsey study: http://www.datacenterknowledge.com/archives/2009/04/15/mckinsey-data-centers-cheaper-than-cloud/. 2. Comparison of IBM System x3850 X5 + MAX5 with total 96 DIMMs x 16 GB for total 1.5 TB of memory vs. IBM System x3850 M2 with 32 DIMMs x 8 GB = 256 GB. Comparison of processor-based licensing fees on current Generation 4 processor systems with 64 DIMMs vs. the IBM System x3690 + MAX5. IBM eXFlash technology would eliminate the need for a client to purchase two entry-level servers and 80 JBODs to support a 240,000 IOPs database environment, saving up to 97% in server and storage acquisition costs.

C O N T E N T S

COVER STORY

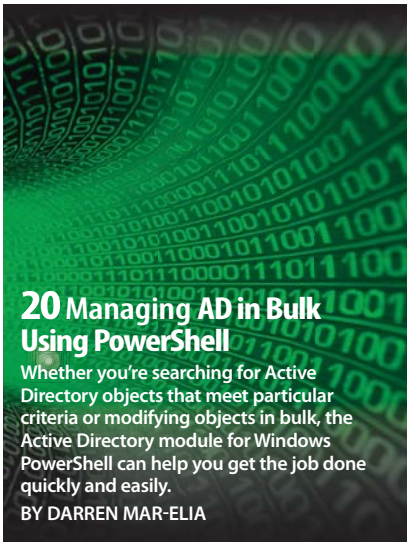

FEATURES

SOLUTIONS PLUS

INTERACT

IN EVERY ISSUE

Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the InstantDoc ID text box on the home page.

Windows IT Pro

A PENTON PUBLICATION

NOVEMBER 2010

VOLUME_16

NO_11

COLUMNS

INFORMATION AVAILABILITY

Virtual · Physical · Cloud

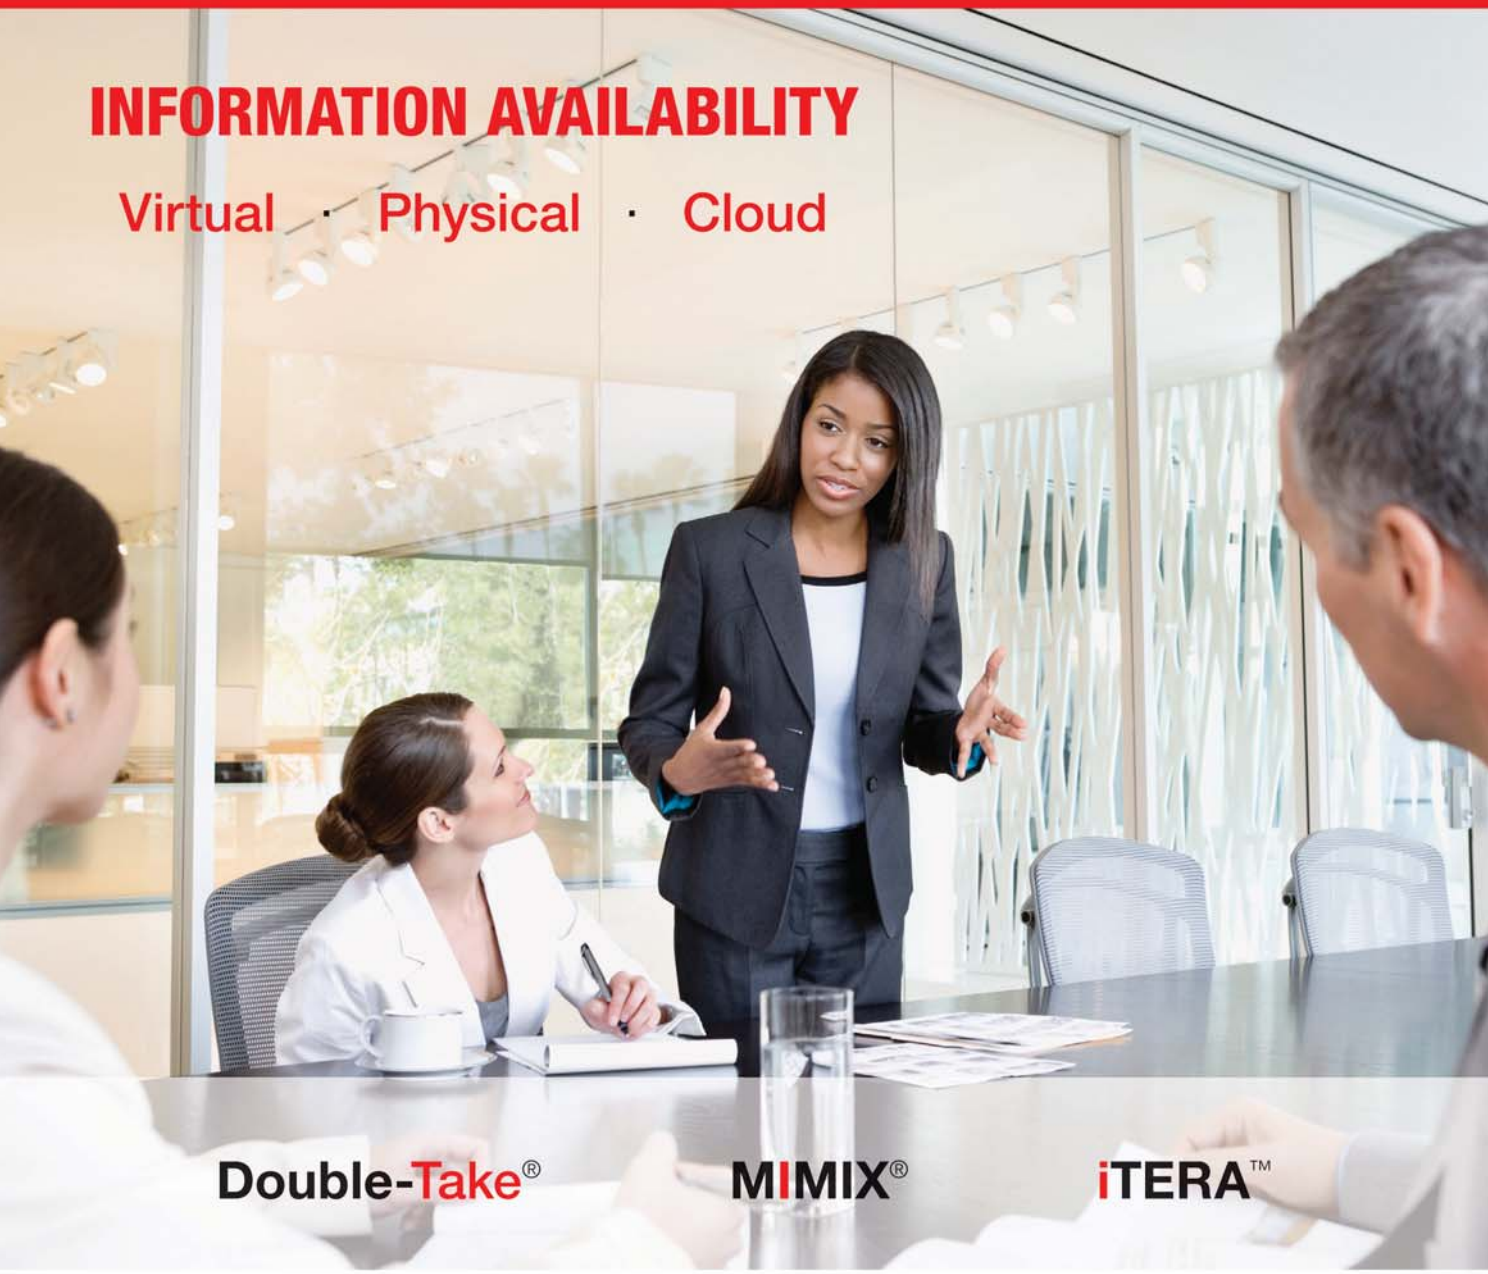

Double-Take® MIMIX® iTERA™

Leaders Have Vision. Vision Has Leaders.

Vision Solutions offers a best-in-class portfolio of high availability and disaster recovery software for Windows®, Linux®, IBM Power Systems™ and Cloud Computing.

The New Vision Solutions features a greater choice of technologies and one of the world's most innovative R&D, service and global customer support teams for availability software.

Take a closer look at visionsolutions.com or call 800-957-4511.

VISION® SOLUTIONS

Easy. Affordable. Innovative. *Leaders Have Vision.*

C O N T E N T S

PRODUCTS

Windows IT Pro

EDITORIAL

Editorial and Custom Strategy Director
Michele Crockett mcrockett@windowsitpro.com

Executive Editor, IT Group
Amy Eisenberg amy@windowsitpro.com

Senior Technical Director
Michael Otey motey@windowsitpro.com

Technical Director, IT Group
Sean Deuby sdeuby@windowsitpro.com

Senior Technical Analyst
Paul Thurrott news@windowsitpro.com

Industry News Analyst
Jeff James jjames@windowsitpro.com

Custom Group Editorial Director
Dave Bernard dbernard@windowsitpro.com

Web and Developer Strategic Editor
Anne Grubb agrubb@windowsitpro.com

Systems Management
Karen Bemowski kbemowski@windowsitpro.com
Caroline Marwitz cmarwitz@windowsitpro.com
Zac Wiggy zwiggy@windowsitpro.com

Messaging , Mobility, SharePoint, and Office
Brian Keith Winstead bwinstead@windowsitpro.com

Networking and Hardware
Jason Bovberg jbovberg@windowsitpro.com

Security
Lavon Peters lpeters@windowsitpro.com

SQL Server
Megan Bearly Keller mkeller@windowsitpro.com
Sheila Molnar smolnar@windowsitpro.com

Editorial Web Architect
Brian Reinholz breinholz@windowsitpro.com

IT Media Group Editors
Linda Harty, Chris Maxcer, Rita-Lyn Sanders

CONTRIBUTORS

SharePoint and Office Community Editor
Dan Holme danh@intelliem.com

Senior Contributing Editors
David Chernicoff david@windowsitpro.com
Mark Joseph Edwards mje@windowsitpro.com
Kathy Ivens kivens@windowsitpro.com
Mark Minasi mark@minasi.com
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com

Contributing Editors
Alex K. Angelopoulos aka@mvps.org
Sean Deuby sdeuby@windowsitpro.com
Michael Dragone mike@mikerochip.com
Jeff Fellinge jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarelia@windowsitpro.com
Tony Redmond 12knocksinna@gmail.com
Ed Roth eroth@windowsitpro.com
Eric B. Rux ericbrux@whshelp.com
John Savill john@savilltech.com
William Sheldon bsheldon@interknowlogy.com
Randy Franklin Smith rsmith@montereytechgroup.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

ART & PRODUCTION

Production Director
Linda Kirchgesler linda@windowsitpro.com

Senior Graphic Designer
Matt Wiebe matt.wiebe@penton.com

ADVERTISING SALES

Publisher
Peg Miller pmiller@windowsitpro.com

Director, International and Agency Services
Don Knox don.knox@penton.com

Business Development Director
Kerry Gates kerry.gates@penton.com

EMEA Managing Director
Irene Clapham irene.clapham@penton.com

Director of IT Strategy and Partner Alliances
Birdie J. Ghiglione birdie.ghiglione@penton.com
619-442-4064

Online Sales and Marketing Manager
Dina Baird Dina.Baird@penton.com

Key Account Director
Chrissy Ferraro christina.ferraro@penton.com
970-203-2883

Account Executives
Barbara Ritter barbara.ritter@penton.com
858-367-8058
Cass Schulz cassandra.schulz@penton.com
858-357-7649

Client Project Managers
Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

Ad Production Supervisor
Glenda Vaught glenda.vaught@penton.com

MARKETING & CIRCULATION

Customer Service service@windowsitpro.com

IT Group Audience Development Director
Marie Evans marie.evans@penton.com

Marketing Director
Sandy Lang sandy.lang@penton.com

CORPORATE

PENTON

Chief Executive Officer
Sharon Rowlands Sharon.Rowlands@penton.com

Chief Financial Officer/Executive Vice President
Nicola Allais Nicola.Allais@penton.com

TECHNOLOGY GROUP

Senior Vice President, Technology Media Group
Kim Paulsen kpaulsen@windowsitpro.com

WRITING FOR *WINDOWS IT PRO*
Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

LIST RENTALS
Contact MeritDirect, 333 Westchester Avenue, White Plains, NY or www.meritdirect.com/penton.

REPRINTS
Diane Madzelonka, Diane.madzelonka@penton.com, 216-931-9268, 888-858-8851

WINDOWS IT PRO VIP is

Educational—with FREE eLearning courses and eBooks available 24×7

Deep—housing over 41,000 articles on DVD and online, some exclusively for VIP members

Broad—solutions, tips, and tricks for any Windows or SQL Server issue that can stump you

In fact, Windows IT Pro VIP delivers more than **$1,000 of resources and expertise for just $199 a year.**

HOW WINDOWS IT PRO VIP BEATS A SEARCH ENGINE		
	Windows IT Pro VIP Delivers:	**Search Engines Deliver:**
Reliability	Road-tested advice from experts who put their reputation on the line	Well-meaning but potentially harmful tips in the latest Wikipedia entry
Speed	The answers you need in seconds searching by keyword, topic, or publication	Lost time spent perusing sites that have mastered search engine rankings but not the art of Active Directory or patch management
Impartiality	Authors and experts who challenge the Microsoft party line and influence industry change	Conventional wisdom touted by industry insiders afraid to tell it like it is

Order Online Now at windowsitpro.com/go/vip

Crockett

"The promise that cloud computing holds has everyone scrambling to find their rightful place in the cloud."

Getting Smart About Cloud Computing

Sort through the hype for the solution that works for you

Want to be the smartest person in your IT organization—or in the whole company, for that matter? Battle your way through the confusion, hype, proof points, and facts about cloud computing to reach the truth about how your company should structure computing services. It's at least as tricky a problem as ending the stubbornly lingering recession.

Windows IT Pro industry analyst Paul Thurrott recently called cloud computing "an ill-defined and rarely understood technology if there ever was one" ("Revolution, Not Evolution," InstantDoc ID 125939). But the promise that cloud computing holds—one of dynamically scalable resources, a power-saving virtualized environment, and pay-as-you-go pricing—is so compelling that everyone is scrambling to find their rightful place in the cloud.

Amidst this headlong rush are some IT pros who are keeping their cool by focusing on technology that delivers real benefits now. Gary Magnuson, an IT specialist with Kroll Factual Data in Loveland, Colorado, made his department's decision to dismiss cloud computing—at least the hosted services version—sound simple. Magnuson, a longtime subscriber to *Windows IT Pro* who joined the staff's recent strategy sessions as a guest panelist, said his team rejected cloud-based services for two reasons: no apparent cost savings and potential security risks.

"We're such a large virtualized machine environment right now, cloud services would not be a cost savings for us, and we would have to worry about security," Magnuson said.

Magnuson, whose team maintains the organization's servers, isn't kidding when he talks about a large virtualized environment. His team recently shifted 800 physical servers to VMs, and in one case he has 140 VMs running on one physical box. Although this move required investment in new hardware to run Microsoft Hyper-V Server 2008 R2, Magnuson said his team looked at the long-term cost savings.

"It costs you in one way, but saves you in the other," Magnuson said. "Also, if you have some servers that are running toward the end of life, instead of buying a new server you can convert it to a VM. Besides getting rid of the physical box, virtualization saves you on power. You're saving the purchase of the physical box itself, physical memory for that box, repairs on the physical box, and power."

For Magnuson's team, the security risks are simply a matter of control. "If applications are on the other side of the firewall, they have to be much more secure," he said. "And sometimes that can get interesting. If a mistake is made in the cloud, you open yourself up to security risks. On the inside of the organization, it's easy to check the firewall. It would be easier for us to host our own private cloud service than to use a public cloud simply because of the security and ease of use. We have a lot of expertise here, and it's easier for us to manage security from inside the organization."

Magnuson said that his team jumped into virtualization when Microsoft released Hyper-V with Windows Server 2008, and their knowledge evolved as they tested and retested scenarios in their lab that they thought could make their jobs easier. His team is constantly looking for workloads to virtualize and improvements to their virtualization environment.

Although Magnuson is a champion of virtualization, he doesn't rule out cloud services for his server team at some point in the future. But in the meantime, he clearly feels that he and his colleagues have the knowledge and expertise to deliver the performance, availability, and security his company needs with a well-tested virtualized environment. He ran through a quick list of factors that businesses need to consider when evaluating cloud services: Would a private cloud service give you better security? Would your business really benefit from the pricing model of public cloud offerings, or would a virtualized environment give you the best cost savings? Are you more confident in a public cloud service's security measures or in your in-house network security expertise? Do you have the resources to manage in-house IT staff? Do you know what resources you'll need to manage outsourced services?

The questions about cloud computing are many, and the answers aren't easy. For a primer on cloud computing, check out senior technical director Michael Otey's article "The Rise of Cloud Computing" (InstantDoc ID 103674). Also, check out the virtualization and cloud computing sessions at Penton Media's Windows Connections conference November 1–4 at Mandalay Bay in Las Vegas (www.devconnections.com). You'll be able to ask our speakers—including Mark Minasi, Sean Deuby, and Mike Danseglio—about the best solutions for your company. The conversation will make us all a little bit smarter about cloud computing.

InstantDoc ID 126037

MICHELE CROCKETT (michele.crockett@penton.com) helped launch *SQL Server Magazine* in 1999, has held various business and editorial roles within Penton Media, and is currently editorial and custom strategy director of *Windows IT Pro, SQL Server Magazine*, and *System iNEWS*.

- Secure NDES
- EFS and Kon-Boot
- Windows 7/XP Dual-Boot
- Cloud Opposition

LETTERS@WINDOWSITPRO.COM

Using Microsoft NDES Beyond Routers and Firewalls

I've just read Russell Smith's "Setting Up Network Device Enrollment Service" (August 2010, InstantDoc 125385). This information comes at a great time for us. We just implemented a new enterprise certificate authority (CA) and are starting to get requests for issuing certificates to cell phones and Linux machines.

Do you have any experience or know of any success stories about folks using Microsoft NDES for other devices outside of routers and firewalls? Can I assume that if the device supports Simple Certificate Enrollment Protocol (SCEP), we should be good to go? Many thanks. I appreciate the great article.

—Thomas Harder

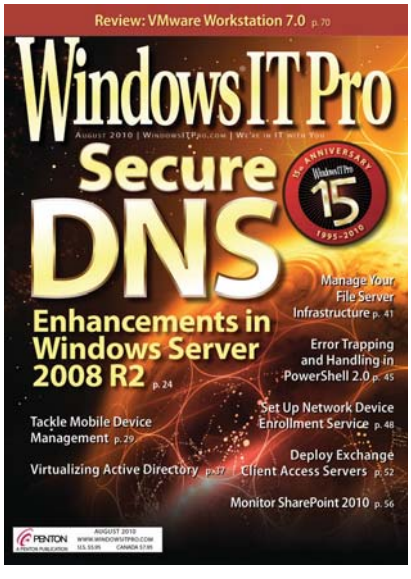

I'm glad you found the article useful! Actually, I wouldn't assume that you're good to go if that device supports SCEP. Different SCEP implementations might not be compatible with one another, so you should test each device that you want to issue a certificate to. Windows Mobile has its own mechanism for deploying certificates (i.e., ActiveSync or the Windows Mobile Device Centre in Windows Vista and later). Apple iPhone (iOS 3.0 or later) is the only mobile device I know of that supports SCEP natively, but it has known issues.

Please let me know if you have success with your SCEP deployment; your story might be something we can feature in a future edition of Windows IT Pro.

—Russell Smith

Dual-Booting Windows 7 and Windows XP

Regarding Michael Otey's "Upgrading from Windows XP to Windows 7" (January 2010, InstantDoc ID 103144)—as well as Robert Schor's comment in the Letters column in the March 2010 issue (InstantDoc ID 103507)—I think both methods are a bit messy. In my environment, I have used the following method, which is simple and works perfectly.

1. Create three partitions—The first two partitions will be the OS partitions, and the third will be the data. Adjust size accordingly. Note that XP will be installed first. You can use the (limited) partition manager when you install XP to set up the three partitions. All three partitions are created and should be formatted. This way, the CD/DVD drives won't change between OSs.
2. Install XP—The middle partition will be the XP installation. When you're done, the C drive will be blank (except for the NT boot files), the D drive will have XP, the E drive will be your data drive, and the F drive will be the DVD drive.
3. Move your My Documents folder to the E drive—This step requires a bit of tinkering. I also move the Favorites folder into My Documents. That way, they'll get backed up, and both OSs can access the same My Documents and Favorites.
4. Install Windows 7 (or Windows Vista, if you wish)—Note that if you're using the 64-bit OS, you must boot off the DVD because you can't (normally) start the 64-bit installation from a 32-bit OS. When prompted, don't upgrade, and aim the installation to the C drive. When you're done, the C drive will be Windows 7, the D drive will have XP, the E drive will be your data drive, and the F drive will be the DVD drive.
5. Move My Documents into the same folders where XP was moved—Do the same for Favorites. In Windows 7, it's as easy as right-clicking the folder in your user file and changing the location.

Now, both OSs will have the same drive sequence. I've encountered only two problems with this method. First, every time you boot into XP, it will kill the Windows 7 System Restore points. To correct the situation, you either "hide" Windows 7 from XP or enable BitLocker in Windows 7. For more information, see the Microsoft article "No restore points are available when you use Windows Vista, Windows 7 or Windows

EFS and Kon-Boot

In his FAQ, "How do I stop tools like Kon-Boot?" (InstantDoc ID 125801), John Savill suggests that one of the mitigations for Kon-Boot is to use Encrypting File System (EFS). But if you have access to local files using cached credentials, wouldn't you also have access to the EFS encrypted files belonging to this user?

—Robin Penny

EFS doesn't really help with the specific Kon-Boot attack because of the cached credentials in use. However, it does help with other types of attacks that focus more on the file system. The FAQ in question refers to "tools like Kon-Boot," so EFS is still a great extra layer of defense. We always want defense in depth where possible. But, for Kon-Boot specifically, BitLocker is king.

—John Savill

Windows IT Pro welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

Server 2008-based operating systems in a dual-boot configuration together with an earlier Windows operating system" (support.microsoft.com/kb/926185).

Also, when you hide Windows 7 from XP, you can't access the XP boot .ini file. However, you don't really need it because both OSs use the Windows 7 boot manager.

—Ed Braiter

A Reader Who Is Highly Opposed to the Cloud

In his Exchange and Outlook blog post, "IT Pros Resist Moving Messaging to the Cloud," (www.windowsitpro.com/blogs/exchangeandoutlook.aspx), B. K. Winstead asks, "It also makes me wonder why, if the IT pros are still so set against cloud computing, the industry analysts and tech speculators continue to tout this as the next big thing?"

Frankly, I find this to be a common occurrence. Although I want to read about the latest technology and its potential benefits, I'm sometimes annoyed when articles slant toward the mindset of "This is the best thing ever, everybody else is doing it, so you should too!" All these new technologies—the cloud, VDI, host VoIP—offer exciting new features and might well be the way of the future, but they're new and they still have drawbacks. Sometimes they lack features we take for granted in our legacy systems.

So, although they might be a fit for some companies, they're not a fit for many—at least not yet. Maybe in 10 years, when it's matured, cloud-based email will be the norm. Until then, journalists and analysts need to remember that most of us have entrenched legacy systems that will be slowly phased out as new technologies mature.

—Richard Van Alstine

Thanks for the response. I'm probably as guilty as anyone of making some of these technologies sound too exciting to readers, but I usually try to look for the loopholes at the same time—which is part of the reason I posted the poll question about hosted Exchange in the first place.

On the one hand, I've long sensed a disconnect between what Microsoft and other vendors are saying about hosted messaging and what's really happening in the field. On the other hand, I'm not sure it will take as long as 10 years for the cloud world to mature enough to gain general acceptance; the technologies just improve too quickly these days.

That doesn't necessarily mean that all companies will find it feasible to move all their systems to the cloud, either. As you say, there are many legacy systems in place, and you're not going to just abandon something that's working.

In my blog, I've recently posted a short piece that's a slight expansion of what I originally wrote: I'd like to encourage you and other readers to post comments there—it's important to continue this discussion and let the IT pro voice be heard!

—B. K. Winstead

InstantDoc ID 125996

Instant Poll Results:
Mac Support

Do you support Macs in your environment?

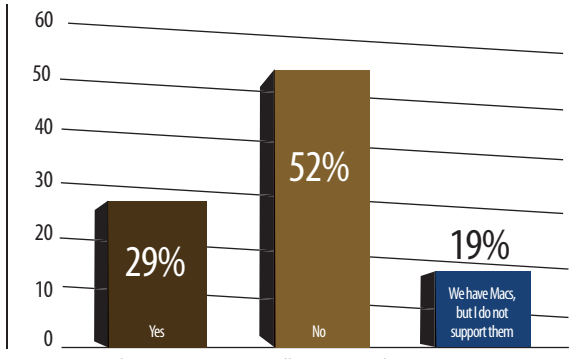

Source: *Windows IT Pro* Instant Poll, www.windowsitpro.com, September 2010

Thurrott

"Under the hood, IE 9 takes advantage of the underlying capabilities of your PC to provide better performance than is possible with other browsers. Text, graphics, and digital media rendering are now hardware accelerated."

What You Need to Know About Internet Explorer 9 Public Beta

An important new software release is coming that warrants your complete attention. Microsoft's next browser, Internet Explorer (IE) 9, is now available in a public beta version. Here's what you need to know.

IE 9 Public Beta

After years of waffling, Microsoft has finally embraced the standards-based web, and it has done so in a big way with IE 9. The software giant's next web browser takes on the same high-level ideals as such products as Windows 7 and Windows Phone—that is, get out of the way visually so that content can take center stage— and it does so in a way that's both attractive and useful. IE 9 really pushes compatibility and performance boundaries while retaining the browser's deployment and management advantages.

User-Experience Changes

The public beta of IE 9, released in mid-September 2010, offers a first look at the browser's streamlined new UI. This follows a series of essentially UI-less developer-oriented public preview releases during which time Microsoft honed the browser's standards-based and hardware accelerated rendering capabilities. See my SuperSite for Windows IE 9 Beta review to check out the screen shots (www.winsupersite.com/live/ie9_beta.asp).

You'll see that IE 9 is clean and minimalistic. Where IE 8 was busy with UI "chrome" (features such as the Favorites Bar and web slices that were arguably useful, but took up lots of onscreen real estate), IE 9 offers up a basic UI that takes up less space onscreen than the UI of any competing browser (let alone previous IE versions).

This is all by design. According to Microsoft, the central design mantra for IE 9 was that the browser must get out of the way, visually. People care about the sites, not the browser, I was told, in the same way that they care about applications, not the Windows OS on which they run.

Speaking of which, one of IE 9's most intriguing usage changes is that the browser promotes websites into first-class Windows citizens, giving them many of the same capabilities as applications, especially under Windows 7. You can now pin website shortcuts to the Windows 7 taskbar (and Start Menu) just as you can application shortcuts. The effect is interesting and immediately logical, since most users access a combination of applications and websites. Putting links to both side by side is natural and intuitive.

Websites that aren't yet ready for IE 9 will display a regular static icon based on the graphic the sites already use. But websites can also be easily modified to take advantage of specific IE 9 and Windows 7 features, including Jump Lists, hover effects, and even pop-up media players (for sites like Pandora). Sites that need to provide notifications—such as email services—can do so via a badge on the site's taskbar icon. And IE uses Aero Snap in Windows 7 to provide drag and drop "snapping" of web pages on the screen so you can view pages side-by-side just as you do with individual applications.

Like Google Chrome, when you do pin a website to the taskbar (or Start Menu), IE creates a specialized version of the IE window designed for that site. The navigational controls and other UI elements are automatically colored to match the design of the site, and the browser's home button is replaced by a site-specific home button so that the site is always "home" for that window. In short, pinned sites are treated like individual applications. It's a neat capability.

Like Chrome, IE 9 dispenses with a separate Search box and integrates search functionality into the address bar, which is now called the One Box. This single UI element can be used to navigate to specific sites, search via the configured search engine, switch search engines, and access browser history and favorites. It won't transmit any of your keystrokes to your configured search engine on the fly unless you okay that behavior. And IE 9 will keep your searches private by default.

The notification bar that debuted back in Windows XP SP 2 has proven so popular that virtually all browser makers have copied the feature. But with IE 9, Microsoft is walking away from this design and offering a new, even less intrusive notification bar that pops up from the bottom of the browser window. And unlike its predecessor, it's not modal, and won't prevent you from browsing along as you ignore it. (There are a few security-oriented exceptions to this rule.)

The new notification bar is also easier to understand, with simple, clear language. When a website hangs, you'll see a warning message that's very similar to what you see when an application crashes.

Tabs get a big boost with IE 9. They can be rearranged and torn off of individual windows and, as noted previously, "snapped" to the screen sides and pinned to the taskbar. There's also a clean-looking New Tab page that simply shows you the sites you've

visited most frequently. There's also a link for InPrivate Browsing, which debuted in IE 8.

From a performance standpoint, IE 9 is a mixed bag in this initial beta release. On a clean install, it's a speed demon, starting up quickly, launching new tabs in a flash, and rendering sites as fast as any other browser. However, most users won't be getting IE 9 as part of a clean install of the OS. Instead, they'll be upgrading from their previous browser. And that's where the problems start, because IE 9 uses the same add-on model as its predecessor, and it loads—and works with—whatever add-ons you've already configured, knowingly or not.

In such cases, IE 9 does throw out one bone: It provides a notification pop-up asking if you'd like to disable add-ons to speed performance. And if you choose to view this UI, you'll see a useful interface that lists your installed add-ons in order, with the slowest performing ones first. A graphical scale explains how long each takes to load, too, so you can gauge the effect of disabling them.

Unfortunately, the Manage Add-ons UI hasn't changed at all. Neither has Internet Options, which dates back over a decade and is overdue for a remake. Microsoft told me it looked into doing so but wanted to focus on the parts of the browser that affect users every day. Fair enough.

Also new in IE 9 is a long-awaited download manager. It integrates with the browser's excellent SmartScreen protection system.

Management and Deployment Features

IE 9 can be deployed through the IE Administration Kit (IEAK) 9, or via Windows Server Update Services (WSUS), or System Center Configuration Manager 2007, and it can be customized. Admins can also slipstream IE 9 into their Windows Vista and Windows 7 install images.

IE 9 also features almost 1,500 Group Policy settings, including many that are new to this release and relate to new IE 9 functionality. For example, you can disable add-on performance notifications, enable newly installed add-ons automatically, or prevent the user from reconfiguring One Box search. For those environments that don't want to install IE 9, or wish to delay the update from Windows Update, Microsoft will again offer a Blocker toolkit.

Under the Hood

Microsoft started pushing its modern and powerful IE 9 site-rendering capabilities back in March 2010 when the first platform preview debuted, but now that the company has shipped a usable version of the browser, we can see how it works in real-world conditions. The prognosis is very positive: In tests over several weeks, I only ran into intermittent issues with a handful of sites, which speaks well to IE 9's compatibility in even this early test version.

Under the hood, IE 9 takes advantage of the underlying capabilities of your PC to provide better performance than is possible with other browsers. Text, graphics, and digital media rendering is now hardware accelerated, as it is in native Windows applications, providing a new benchmark for browser performance.

IE 9 also features almost 1,500 Group Policy settings, including many that are new to this release and relate to new IE 9 functionality.

Microsoft notes that only IE 9 offers true, across-the-board hardware acceleration that fully utilizes the underlying Windows capabilities. Other browsers, constrained by their cross-platform needs, can offer only partial acceleration, and that's using inefficient intermediate cross-platform code layers.

IE 9 is the most standards-compliant browser Microsoft has ever created, though it still lags behind Chrome and Safari. It supports many HTML 5 features, including video and audio playback that doesn't require a plug-in, full Cascading Style Sheet (CSS) 2.1 and partial CSS 3 support, ECMAScript 5 (the successor to JavaScript) support, and more.

IE 9's new scripting engine, called Chakra, utilizes multiple CPU cores to interpret, compile, and run code much more quickly than is possible with IE 8, offering performance similar to that seen in the fastest browsers available today. HTML and CSS rendering, scripting, formatting, layout, and other activities all contribute to the real-world rendering of sites, as does IE 9's new hardware acceleration capabilities, which offload tasks from the CPU.

What Microsoft Learned With IE 9

One side effect of the IE 9 development schedule is that the software giant was able to quickly evolve the product based on feedback. Microsoft says the IE 9 development process was a useful experience that will guide subsequent releases. When you consider the speed at which its competitors move, it's hard not to imagine Microsoft at least cutting the difference between the schedules of those products and its previously stately progress.

The Schedule to Come

Microsoft tells me that it will issue a release candidate version of IE 9 in the months ahead, followed by the final shipping version of the product. Both of these upcoming milestones will be timed based on the feedback Microsoft receives on the preceding release, so it's unclear if the RC or final version of IE 9 will ship in late 2010 or early 2011. I'm thinking early 2011 for the final version.

Final Thoughts

So far, I'm deeply impressed by the work Microsoft has done to make IE 9 usable, fast, and standards compliant. The software giant has traditionally held back on the latter quality in particular, due to compatibility concerns. Whether this is still an issue remains to be seen, but I could see some legacy intranets causing issues, so testing the new browser early will be key. At the very least, IE 9's legacy compatibility modes should ease some of that burden.

InstantDoc ID 125989

PAUL THURROTT (thurrott@windowsitpro.com) is the news editor for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

Minasi

"Now that you've met Imagex /capture and Imagex /apply, you can start creating and working with WIM files."

Deploying and Error Checking with ImageX

Two recommended options—/verify and /check—bring peace of mind to your apply and capture operations

In "ImageX Provides Disk Imaging on a Budget" (InstantDoc ID 125743), I introduced ImageX, a Microsoft tool that—similarly to Symantec's Ghost—lets you easily clone an entire drive from one system and distribute it to as many other PCs as you like. I explained where to find ImageX and how to use it to capture the C drive of a working Windows box (which you've already prepared with the Sysprep utility), by using a command such as

```
imagex /capture C: G:\baseimage.wim "Base Win 7 image"
```

In that example, I've captured the entirety of a computer's C drive onto a file on G named baseimage.wim. Now that you have this image file, what can you do with it?

Clearly, your first use of baseimage.wim would be to deploy it to a PC. That process involves four steps: Boot the target PC with Windows Preinstallation Environment (WinPE), wipe any existing partitions, create and format a new partition, and use ImageX to apply baseimage.wim to the new partition. (Note that I'm assuming the system you created was formatted as one large partition—something I'll cover in greater detail in the near future.) I discussed WinPE in the previous article, and in the past year I've covered Diskpart—a command that lets you create, destroy, and format partitions—but in case you've forgotten, the following commands will get you ready to deploy on most systems:

```
diskpart
select disk 0
clean
create partition primary
format fs=ntfs quick
assign letter=C
exit
```

Now, your target system's hard drive is formatted as one large drive named C. (Any drive letter will do, because you're merely in WinPE, and once you've deployed the image and reboot, Windows will automatically re-letter the drive to C.) Next, you'll need to get connected to whatever storage contains baseimage.wim—whether it's on a USB stick, an external drive, or a network share. Let's assume that you've mapped the storage as drive P. To apply baseimage.wim to the new C drive, type

```
imagex /apply P:\baseimage.wim 1 C:
```

(The 1 indicates that you want to use the first image in the baseimage.wim file, and I'll cover that in more detail next month.) After ImageX does its work, just reboot the newly imaged system. Depending on how you prepared your prototype system with Sysprep, either you'll have to respond to the usual *Name the computer, Create a user,* and *Set Windows Update settings* prompts or—if you created an automated setup script for Sysprep—you'll see the system start up, ready for you to log on and get to work.

Now that you've met Imagex /capture and Imagex /apply, you can start creating and working with WIM files. But ImageX is capable of much more, so let's meet the /verify and /check options.

The /verify option ensures that either a capture or apply operation didn't accidentally drop any bits along the way. WIM files are pretty large, and although network and disk read/write operations already contain built-in checks, just one misplaced byte can render the image or system useless. The /verify option helps prevent that from happening. After either applying or capturing a windows image, ImageX compares the source data and its copy, searching for (and correcting) differences. ImageX enables the /verify option by default whenever applying or capturing over a network connection, but it doesn't do that when working with local external storage, so it's never a bad idea to add the /verify option. (It will, of course, slow things down a bit.)

The /check option has a similar purpose but a slightly different and perhaps more efficient approach. If you include /check when capturing an image, ImageX creates a hash of every 10MB chunk of ImageX data, then embeds those hashes in the resulting WIM. Including /check on an apply operation causes ImageX to check the applied image by re-computing hashes on that image, then comparing those computed hashes with the embedded ones. Thus, /check certifies not only that a WIM wasn't mis-copied during the apply operation but also that a WIM hasn't become corrupted while stored. Why use /verify, then? Why not just always use /check? The answer is that /check can't do its job on an apply operation if no one specified /check on the capture. If the WIM lacks embedded hashes, /check can't check the hashes. The /verify option, in contrast, doesn't need hashes, so it can always help.

In any case, Microsoft recommends using both options, so if you've got the time, don't hesitate to verify and check!

InstantDoc ID 125970

MARK MINASI (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of more than 25 books, including *Mastering Windows Server 2008 R2* (Sybex).

NetWrix
Systems Management and Compliance

Yet Another 10 Free Tools for System Administrators

Audit Active Directory and file servers, detect inactive users, block USB devices, and more – for free

The following freeware tools by Windows IT Pro Community Choice Awards finalist NetWrix Corporation can save you a lot of time and make your network more efficient – at absolutely no cost. Some of these tools have advanced commercial versions with additional features, but none of them will expire and stop working when you urgently need them.

1. Active Directory Change Reporter (Windows IT Pro Sep'09: InstantDoc ID 102446, Windows IT Pro Jan'09: InstantDoc ID 100593, TechTarget: www.tinyurl.com/322zf9x) — This is a simple auditing tool to keep tabs on what's going on inside Active Directory. The tool tracks changes to users, groups, OUs, and other types of AD objects, and sends summary reports with full lists of what was changed and how it was changed. In addition, it has a nice "rollback" feature that helps rollback unwanted changes (including deletions) very quickly. Download link: www.tinyurl.com/24sqv7c

2. USB Blocker (Windows IT Pro Nov'09: InstantDoc ID 102860) — Users bring tons of consumer devices: flash drives, MP3 players, cell phones, etc., into the office and this aptly-named tool can block them with a couple of mouse clicks to prevent the spread of a virus and to restrict the take-out of confidential information. The product is integrated with Active Directory and is very easy to use. Download link: www.tinyurl.com/24ls5zx

3. Password Expiration Notifier (Redmond Magazine Feb'09, 4sysops: www.tinyurl.com/2ubnqnu) — This tool will automatically remind users to change passwords before they expire to keep you safe from password reset calls. It works nicely for users who don't log on interactively and, thus, never receive standard password change reminders at log on time (e.g., VPN and OWA users). Download: www.tinyurl.com/258elbh

4. Inactive Users Tracker (MS TechNet Magazine May'08: www.tinyurl.com/2u2wm9g) — This feature tracks down inactive user accounts (e.g., terminated employees) so you can easily disable them, or even remove them entirely, to eliminate potential security holes. The tool sends reports on a regular schedule, showing what accounts have been inactive for a configurable period of time (e.g., 2 months). Download link: www.tinyurl.com/2fycw5b

5. File Server Change Reporter (4sysops.com: www.tinyurl.com/2wz3pgh) — This tool enhances the line of auditing tools; this one for file servers. File Server Change Reporter detects changes in files, folders, permissions, tracks deleted, and newly-created files, and sends daily summary reports. This is a very useful tool to detect mistakenly-deleted files and recover from backup or to see if someone changes some important files. Download link: www.tinyurl.com/2eb2d6u

6. Active Directory Object Restore Wizard (4sysops.com: www.tinyurl.com/3xbmv2c) — This tool can save the day if someone accidentally (or intentionally) deleted a bunch of Active Directory objects. It provides granular object-level and even attribute-level restore capabilities to quickly rollback unwanted changes (e.g., mistakenly deleted users, modified group memberships, etc). Download link: www.tinyurl.com/23gflz3

7. VMware Change Reporter (TechTarget/SearchVirtualDesktop: www.tinyurl.com/37vu68s) — If you don't know what is being changed by your colleagues in the VMware infrastructure, it's very easy to get lost and miss changes that can affect the things for which you are responsible. This tool tracks and reports configuration changes in VMware Virtual Center settings and permissions. Download link: www.tinyurl.com/2cujnrt

8. Windows Service Monitor (WindowsReference.com: www.tinyurl.com/2v6syb2) —This very simple monitoring tool alerts you when some Windows service accidentally stops on one of your servers. The tool also detects services that fail to start at boot time, which sometimes happens, for example, with Exchange Server. Download link: www.tinyurl.com/28yx29o

9. Bulk Password Reset (reviewed by SoftPedia: www.tinyurl.com/3abzme3) – While most companies have strong password policies for their employees, one critical issue is still neglected: local Administrator passwords on all servers are usually managed in a "set and forget" fashion, sometimes using some "well-known" passwords, opening a major surface for security attacks. The Bulk Password Reset tool quickly resets local account passwords on all servers at once, making them more secure. Download link: www.tinyurl.com/27bmaj7

10. Disk Space Monitor (MS TechNet Magazine Sep'09: www.tinyurl.com/3225fg6) — Even with today's terabyte-large hard drives, server disk space tends to run out quickly and unexpectedly. This simple monitoring tool will send you daily reports regarding all servers that are running low on disk space, below the configurable threshold. Download link: www.tinyurl.com/23kt6lt

Otey

"Hyper-V is making solid inroads into Windows IT infrastructures everywhere."

Tips for Using Hyper-V

Tricks and techniques for managing and getting the best performance from your VMs

Microsoft's Hyper-V virtualization platform has turned out to be a reasonable competitor to VMware's ESX Server. Hyper-V is making solid inroads into Windows IT infrastructures everywhere. To help you make the most of your Hyper-V installations, here are ten tips for managing and getting the best performance from the Hyper-V platform.

10 **Uninstall VM Additions before migrating VMs—**If you're moving virtual machines (VMs) from Microsoft Virtual Server or Virtual PC to Hyper-V, uninstall the existing VM Additions while the VM is still running on the old platform. You can't uninstall them after moving the VM to Hyper-V.

9 **Use Import and Export to manually move VMs between hosts—**Live Migration is a great feature, but not everyone has the infrastructure required to support it. Moving Hyper-V VMs isn't as easy as it was with Virtual Server or Virtual PC. To move a VM, open the Hyper-V Manager, select a VM, then choose the Action, Export option. On the target Hyper-V server, select Import Virtual Machine from the Actions pane.

8 **Enable Windows Backup—**Windows Server's Volume Shadow Copy Service (VSS) lets you back up live applications with no downtime. Although Hyper-V is VSS-aware, VSS compatibility isn't enabled out of the box—it requires adding a registry entry. For details about the process, check out the Microsoft blog "How to enable Windows Server Backup support for the Hyper-V VSS Writer" (bit.ly/bf1a2r).

7 **Make sure you have enough physical network adapters in the host—**During server consolidation, it's easy to overlook the number of physical network adapters you have in the Hyper-V host. Multiple VMs sharing network adapters can easily overload the available bandwidth. Ideally, you should dedicate a NIC for each VM. If that isn't possible, spread the network load across as many physical network adapters as you can.

6 **Manage Hyper-V with PowerShell—**Microsoft System Center Virtual Machine Manager (VMM) has a PowerShell provider you can use to manage Hyper-V. However, most organizations don't use either System Center or VMM. That doesn't mean you need to miss out on PowerShell management. There's a free PowerShell Management Library for Hyper-V available on CodePlex at pshyperv.codeplex.com.

5 **Take advantage of improved dynamic disk performance in Hyper-V R2—**With the original version of Hyper-V that was released with Windows Server 2008, many administrators chose to avoid using dynamic disks because of the performance overhead associated with them. Hyper-V R2, which was released with Server 2008 R2, overhauled the performance of dynamic disks, making them nearly as fast as fixed virtual disks.

4 **Access iSCSI storage from VMs—**Hyper-V's New Virtual Machine Wizard makes it easy to configure access to DAS from a VM. However, you can also connect VMs to storage on iSCSI SANs. To connect a VM running Server 2008 to an iSCSI SAN, use the Administrative Tools, iSCSI Initiator option and point the iSCSI Initiator to your iSCSI SAN.

3 **Boot VMs off iSCSI LUNs—**You can boot VMs off iSCSI LUNs by attaching the LUNs to the Hyper-V parent partition. Next, create the VM without configuring a hard disk. After the VM is created, select the IDE controller and add a hard disk. On the Hard Disk configuration dialog box, select *Physical hard disk*, then select the iSCSI LUN. When you start the VM, it boots off the iSCSI LUN.

2 **Take advantage of processor compatibility—**Live Migration lets you move VMs between different Hyper-V hosts with no end-user downtime. However, the hosts have to have compatible processors: They must be from the same manufacturer and family. If your processors don't match, you can enable processor compatibility by opening the VM settings, clicking the Processor node, then checking the *Migrate to a physical computer with a different processor version* checkbox.

1 **Watch out for virtual DC gotchas—**If you're running a Domain Controller (DC) in a Hyper-V VM, you need to look out for a few gotchas. First, don't use the save state option for the virtual DC because it can cause synchronization problems in the domain. If you need to stop the system, you should select the VM guest's Shut Down option. You shouldn't pause a virtual DC for more than a couple of minutes because this process can disrupt replication. For the same reasons, you shouldn't take or restore snapshots of the VM acting as a DC.

InstantDoc ID 125932

MICHAEL OTEY (motey@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Gerber

"DNS is the bedrock of your AD environment; without name resolution, operations may grind to a halt."

Identify and Troubleshoot DNS Problems

4 expert tips to help DNS work better in your AD environment

When DNS name resolution stops working, your Microsoft Active Directory (AD) environment has problems. DNS is the bedrock of your AD environment, and without name resolution, operations may be interrupted or grind to a halt. Possibly a name resolution issue has built up over time due to a slow migration away from a sound DNS tree hierarchy design. Here are some key tips and troubleshooting tools to help you avoid or resolve DNS problems.

Tip 1: The DNS namespace should reflect a contiguous tree hierarchy

The Internet DNS namespace has a tree hierarchy (by design), and administration of this is delegated to DNS administrators responsible for various branches of the DNS namespace (IETF RFC 1034, RFC 1032). Like their Internet counterparts, intranet DNS admins should follow hierarchical tree design. Whenever a non-contiguous or disjoint namespace is encountered in an intranet environment, complexity will ensue in the form of the addition of conditional forwarders, stub zones, and/or secondary zones.

Take the scenario of a company acquisition, where each company has two separate, independent namespaces that must be consolidated in functionality using Windows Server trusts. A suggested approach is to create a secure VPN between both company environments at the root of each separate, continuous namespace. Use conditional forwarders to direct name queries across the VPN to the namespace on the opposite side of the VPN tunnel, as Figure 1 shows. Any queries that do not meet conditional forwarder criteria would simply be forwarded up to the ISP's DNS servers. DNS servers configured for conditional forwarding located in the root level of one contiguous namespace will forward queries to a specific DNS server located in the other contiguous namespace. A cache of namespace information is amassed on this DNS server, and the need for recursion is decreased.

Instead of using conditional forwarders, some admins choose to implement stub zones on the DNS servers in top-level intranet domains such as domain1.local and domainA.local. Stub zones contain only enough record information to be able to determine the authoritative DNS servers for the subordinate zone and are more of a consideration when zones are not stored in AD. Stub zones are a consideration when a DNS server in a parent zone needs to be kept aware of the authoritative DNS servers in the child zone. Stub zones increase complexity because when a stub zone replies to a query for a name, all authoritative DNS servers in the domain are provided in the DNS response. The goal should be to have a DNS infrastructure design that is functional and straightforward to troubleshoot.

Tip 2: Understand where DNS information is stored

DNS zone data can be stored in the AD information tree or in the file system in c:\%systemroot%\system32\dns. I strongly recommend that you store zone information in AD, then replicate this zone information either to every DNS server in the domain (DomainDNSZones) or possibly in the forest (ForestDNSZones). Storing DNS information on every DNS server in the domain, then forwarding upstream to the parent zone is an optimal choice. DNS forwarding would be set up so that DNS servers in child1.domain1.local and child2.domain1.local forward to DNS servers in parent domain1.local. In the parent domain, there would be delegation to each child domain. (For additional information about DNS zone location, see "Chasing the DNS Zone Location Problem," May 2010, tinyurl.com/2bst9tz.)

Tip 3: Identify whether the DNS problem is a name-registration or name-resolution problem

To resolve a name, the name must be registered in a zone on a DNS server. In a Windows environment, different services register different records. In Windows 7, Windows Vista, Windows Server 2008 R2, and Server 2008, the DNS client service registers A and PTR records. In Windows XP and Windows Server 2003, the DHCP

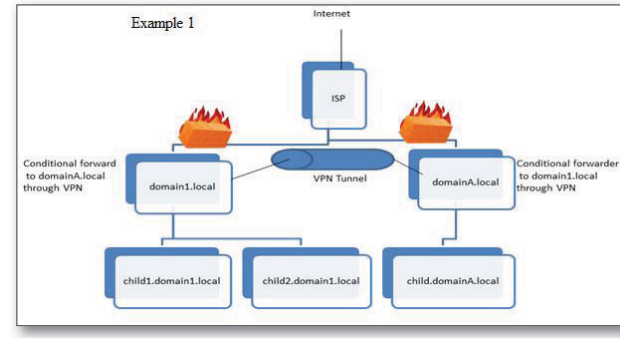

Figure 1: Using conditional forwarding across separate DNS namespaces on an intranet

client service registers A and PTR records. The registration interval is 24 hours, except for when the DHCP server is doing the registering; in this case, the registration should take place when the DHCP client's lease is renewed.

In Server 2008 R2, Server 2008, and Windows 2003, the Netlogon service is responsible for the registration of some additional records. A log of the records registered by the Netlogon service is located at %SystemRoot%\System32\Config\Netlogon.dns. Domain controllers (DCs) dynamically register 15 to 30 SRV records every hour in Server 2008, whereas in Windows 2003 the registration by Netlogon is every 24 hours.

In Server 2008, the Cluster service registers the cluster network name resource when the resource comes online. The record is updated at least once every 24 hours. The setting RegisterAllProvidersIP can be used to determine whether all IP addresses for the network name resource are registered in DNS. (For more information, see the Microsoft article at support.microsoft.com/kb/947048.)

The problem is a DNS registration issue. If DNS records are not present in the DNS console, use ADSI Edit to verify that the records are not simply being displayed in the DNS console GUI or in AD. Verify record existence in AD by following the steps in the article at support.microsoft.com/kb/867464. If the records are not present, install Microsoft's Network Monitor on the machine performing the DNS registration and take a network trace while attempting to register the A, PTR, or SRV records. To initiate A and PTR record registration, issue this command:

```
ipconfig /registerdns
```

For SRV record registration, issue this command:

```
c:\net stop netlogon && net start
  netlogon
```

Stop the network trace and filter on DNS traffic. If no registration traffic is present in the network trace, focus on whether the service responsible for the registration (DHCP client, DNS client, Netlogon, Cluster) is running, and check the event logs. (If you're still stuck at this point, it may be time to call Microsoft Support.)

The problem is a DNS resolution issue. If the technical issue is not related to DNS record registration, change the troubleshooting approach and investigate DNS name resolution. First, ping the Fully Qualified Domain Name (FQDN) of the target and determine success or failure. If the failure is by name and not by IP address, verify that the DNS server settings are properly configured in the TCP/IP properties of the machine initiating the query. Next, start a network trace and clear the resolver cache by issuing this command:

```
c:\ipconfig /flushdns
```

Now ping the target by FQDN (e.g., ping server.domain1.local). Stop the network trace and determine whether there is an outbound DNS query and/or an inbound DNS response. The goal here is to determine whether the issue is with getting a query to the DNS server or if the DNS server gets the query and either doesn't respond or the response fails to reach the DNS query initiator.

Tip 4: Use DNS diagnostic tools

To assist you in troubleshooting DNS issues, make sure you have these tools in your DNS toolkit: DNSLint, DCDiag, and NSlookup.

DNSLint. The DNSLint utility has three functional tests, all of which output results to an HTML report. The test are for "lame delegation", the DNS records required for AD replication to succeed, and verifying a user-defined set of DNS records on multiple DNS servers. Specify /d on the dnslint command to perform the domain name test and provide results that can help in diagnosing lame delegation. Specify /s to indicate the IP address of the DNS server for the DNS server authoritative for the domain. Specify /ad to determine whether the DNS record needing AD forest replication is resolvable. (For more information, see support.microsoft.com/kb/321045.)

DCDiag. You can run the dcdiag command using the option /test:DNS. Test options include a DNS basic test and tests for forwarders and root hints, delegation, DNS dynamic updates, DNS record registration, and Internet name testing.

Test the health of a DC:

```
DCDIAG /TEST:DNS /v /s:<DCNAME>
  /f:<filename.log>
```

Test the health of all forest DCs:

```
DCDIAG /TEST:DNS /f /e /f:<filename.log>
```

Test the DC's ability to register the DC Locator DNS records:

```
DCDIAG /TEST:RegisterInDNS
  /DnsDomain:<FQDN of domain>
  /v /f:<filename.log>
```

(In the previous commands, /v specifies verbose output, /s means run local, /f means direct output to file, and /e means test all servers.)

In Windows 2003 SP2, use the DCDiag utility included with SP2, as described in support.microsoft.com/kb/926027. In Server 2008 and Server 2008 R2, install DCDiag by navigating to Server Manager, Features, Add Features, Remote Server Administration Tools, Role Administration Tools, Select DNS Server Tools, Next, Install.

NSlookup. This is a well-known command for DNS troubleshooting. View NSlookup syntax variations by running NSlookup from a command prompt, then issuing the command *help*. Keep in mind that NSlookup has its own built-in stub resolver in the executable and does not use the OS's resolver.

Healthy DNS, Healthy AD

A Windows AD environment can experience a variety of problems when name resolution fails. Determine whether the problem is localized to a machine, subnet, or network. Next, determine whether the problem is with DNS name registration or with DNS name resolution. Finally, use Microsoft tools when needed, both for troubleshooting and for keeping your DNS environment healthy. (Learn more about these tools in the online version of this article at www.windowsitpro.com, InstantDoc ID 125990.)

InstantDoc ID 125990

BOYD GERBER (boydg@microsoft.com) is a support escalation engineer with Microsoft Networking Escalation. His DNS industry experience is with DNS support and DNS development teams. He has a graduate degree in software engineering from the University of Texas at Austin.

READER TO READER

Tips for Optimizing PowerShell Scripts

Some Windows PowerShell scripts can be useful but consume a great deal of memory, take a long time to complete, or both. Here are some tips on how to improve the performance of such scripts.

Optimize last. Don't try to optimize PowerShell scripts as you write them. You might be optimizing code that either disappears on its own or doesn't have a significant effect on final performance.

Use filtering parameters. PowerShell can consume a lot of resources because some cmdlets are designed to provide immense quantities of data. So, if a cmdlet has filtering parameters (-Filter, -Include, and -Exclude), use them.

If a cmdlet supports the -Filter parameter, you want to use it first. It uses the underlying APIs for an object, which means the code is extremely fast because the filter is applied before the cmdlet creates the objects. The -Include and -Exclude parameters are applied to objects after the cmdlet has already created the objects but before the objects go into the PowerShell pipeline. So, they're slower than -Filter, but they're still faster than filtering after the objects are in the PowerShell pipeline.

Sometimes you should use more than one type of filtering. For example, suppose you're searching for all files on the D drive with the file extension .htm. The -Filter parameter uses the traditional Windows file-system semantics, where *.htm returns all files whose extensions begin with .htm. The Windows APIs implement this filtering, making it extremely fast. However, they're very old and consequently just ignore anything in a file extension beyond the first three characters. So a search only for *.htm with -Filter would also return .html files, for instance. Therefore, for speedy filtering, you should use both -Filter (to cut out the vast bulk of files before loading them) and -Include (to get only .htm files). Callout A in Listing 1 shows what this looks like.

Remember, though, that -Filter uses the underlying APIs, so how fast it works depends on those APIs. Take, for example, the code in callout B. In this case, -Filter works slower because the Get-WmiObject cmdlet uses the WMI Scripting API. It's also slower because the WMI Query Language (WQL) is being used for filtering, so the filtering occurs within WMI.

Reduce resource usage. Performance optimization is about reducing resource usage as well as reducing execution time. Sometimes you can do both. Other times you have to make a choice. For example, suppose you need to do something to each file on the D drive. As callout C shows, you could use the ForEach-Object cmdlet to go through all the file-system objects in the collection. When you use this cmdlet, each object goes through extra packaging work when crossing the pipeline boundary, which slows the code down significantly. However, it doesn't consume much memory because only one object passes through the pipeline at a time.

Alternatively, you could use the iterative foreach loop in callout D. This loop takes less time to run because it avoids pipeline boundaries. However, it collects all the file-system objects before processing them, so it can consume excessive system resources if the collection is large.

The foreach loop is faster but uses more memory than the ForEach-Object cmdlet. So, the foreach loop is generally a better choice if you don't expect to have large data sets.

Throttle CPUs with Sleep. PowerShell code that touches many objects often requires a long time to execute and might not yield processor time willingly. This is less troublesome than it was in the days of single-core CPUs, but it still can cause the system to spend a lot of time waiting for things to happen. If your code consumes a lot of CPU cycles or needs to wait for something to happen, you can use the Start-Sleep cmdlet. By default, Start-Sleep operates in seconds, but you can specify a pause time in milliseconds (ms). Clock resolution is typically no better than 10 to 20 ms, so the shortest sleep time you'll probably want to specify is 20 ms. In addition, you don't need to have a pause every loop iteration. Yielding every few iterations is sufficient to ensure that the current CPU is also available for other work. The loop in callout E uses Start-Sleep to pause for 20 ms every 10 items.

You can combine these tips into a basic plan for optimizing scripts. First, don't worry about optimization until the script is complete. Next, when possible, use -Filter to restrict the number of items read into your script and use -Include and -Exclude for further tweaking. This should reduce both running time and resource use. At that point, if you still have numerous items, consider replacing ForEach-Object pipeline elements with foreach loops to speed up the script. (Note that if you're looking at hundreds of thousands of items, this can cause other performance problems.) Finally, if you find that your script has excess CPU consumption, you can use Start-Sleep in core loops.

—Alex K. Angelopoulos, IT consultant

InstantDoc ID 126046

Listing 1: Sample PowerShell Commands

```
A Get-ChildItem -Path D:\ -Filter *.htm -Include *.htm -Recurse

B Get-WmiObject -Class Win32_Product `
  -Filter 'Vendor LIKE "%Microsoft%"'

  # {. . .} represents the code being run on each file.
C Get-ChildItem -Path D:\ -Recurse | ForEach-Object {. . .}

D foreach($file in (Get-ChildItem -Path D:\ -Recurse)) {. . .}

E $i=0
  Get-ChildItem -Recurse |
   ForEach-Object{
    $i+=1
    if($i%10 -eq 0){sleep -mill 20}
   }
```

SharePointPro CONNECTIONS

Join the SharePoint Expert Community

SharePointPro Connections

provides real-world advice from professionals and peers who share their experience administering and developing in SharePoint.

SharePointPro Connections is the independent voice on SharePoint technology. Expert authors provide our community members with field-tested information to enable content and image management, collaboration, and workflow solutions tailored to business needs.

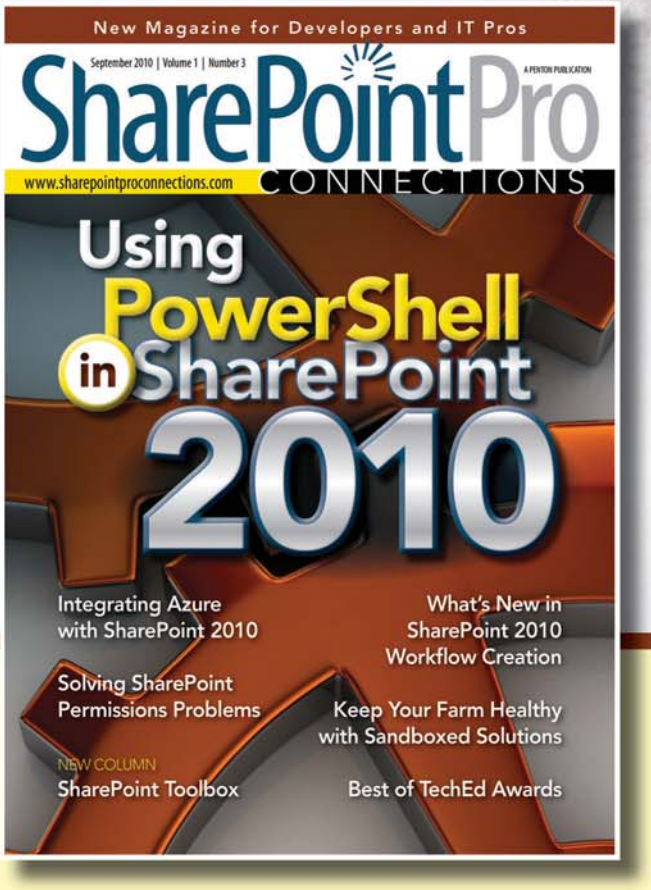

Upcoming topics include:

- Developing Line of Business Connectivity with SharePoint BCS (Business Connectivity Services)
- Planning for SharePoint in the Cloud and on Mobile Devices
- Deep Collaboration with SharePoint Social Media
- Developing and Administering SharePoint Business Intelligence (BI) Solutions

Subscribe FREE to the only magazine dedicated to all things SharePoint.

sharepointproconnections.com/go/SubscribeNow

- PowerShell
- ESX
- Storage
- Hyper-V
- Outlook

ANSWERS TO YOUR QUESTIONS

Q: How can I hide a drive from a user?

A: There may be environments where a drive letter should be hidden from a user because you don't want them using it. For example, you might have a drive where you store a certain type of system cache, or a Citrix Virtual Desktop Infrastructure write-cache.

You can hide drives using Group Policy. Go to User Configuration, Policies, Administrative Templates, Windows Components, Windows Explorer, Hide these specified drives in My Computer. By default, you have options to hide the A, B, C, and D drives. If you want to hide additional drives, you can modify the system.adm file, as detailed at support.microsoft.com/kb/231289.

—John Savill

InstantDoc ID 125965

Q: How can I set the color scheme in Outlook 2007 and Outlook 2010?

A: Office 2007 and Office 2010 both provide the option to change the main color scheme. This can be done from any of the main Office Suite applications, including Outlook. In Outlook 2010, this is accessed through the Office Backstage by selecting File, Options. In the General tab, the Color Scheme lists three options in its drop-down box: Blue, Silver, and Black. If you use Microsoft Expression, you may be accustomed to the default black color scheme and would like to have Office follow that style.

The value for this selection resides in the registry as well, at HKCU\Software\Microsoft\Office\<Office_Version>\Common\, where <Office_Version> is 12.0 for Office 2007 and 14.0 for Office 2010. The Theme DWORD value determines the color scheme for Office 2010:

1. Blue (Default)
2. Silver
3. Black

Changing this value to a number beyond this list didn't change the color scheme from whatever it was set to already. This is one of many settings that can be controlled with Group Policy in a Group Policy Object (GPO) and the Office Customization Tool (OCT).

—William Lefkovics

InstantDoc ID 125895

Q: How do I resolve the error message "The virtual machine is installed VMware Tools and cannot initiate a migration operation"?

A: Sometimes you may attempt to vMotion an ESX 4.x virtual machine (VM) and receive the error message The virtual machine is installed VMware Tools and cannot initiate a migration operation. This error message typically occurs after upgrading a VM's VMware Tools from one version to another.

In this case, VMware is unable to vMotion the VM because the VMware Tools installation has hung. In some situations, this hang situation cannot be cancelled from inside the vSphere client.

There's an alternate way to cancel the VMware Tools upgrade installation within the ESX Service Console using two commands. Prior to using these commands, you'll need to know the full path to the VMX file for the affected VM. With this information, enter the command

```
/usr/bin/vmware-cmd /vmfs/volumes/
  {datastoreName}/{vmFolder}/{vmxFile}
  .vmx getid
```

The result of this first command will be the ID number of the VM. Replace <idNumber> in the command below with this number.

```
/usr/bin/vmware-vim-cmd vmsvc/tools
  .cancelinstall <idNumber>
```

Running this command should cancel the VMware Tools installation and allow you to again vMotion the VM.

—Greg Shields

InstantDoc ID 125976

Q: Are 64-bit and 32-bit PowerShell the same?

A: Nope. For the most part, you won't encounter any differences, but each version of the shell can only load matching snap-ins in some cases, meaning you'll have to be careful to download the correct 64- or 32-bit edition of any add-ons you want to use. Apart from add-on compatibility, there aren't any major functional differences between the two versions.

—Don Jones

InstantDoc ID 125889

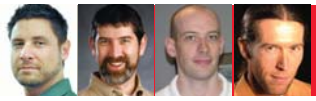

Don Jones | powershell@concentratedtech.com
William Lefkovics | william@mojavemediagroup.com
John Savill | jsavill@windowsitpro.com
Greg Shields | virtualgreg@concentratedtech.com

Q: How can I set Root Memory Reserve for My Hyper-V R2 SP1 parent partition?

A: Microsoft's new dynamic memory feature in Windows Server 2008 R2 SP1 asserts that we IT professionals no longer have to make arbitrary decisions about how much RAM to assign to a virtual machine (VM). Instead, you should let your VMs tell you how much memory they need and let your Hyper-V hosts assign that memory on the fly. It's a neat concept, one that has the potential for dramatically changing how we administer our Windows servers (and desktops in VDI environments).

However, as you can imagine, this level of memory assignment dynamics can create a situation where too much memory is assigned to VMs, leaving nothing for the host's needs. In environments that follow Microsoft best practices, this situation shouldn't happen all that often, because the best practices stipulate that Hyper-V hosts should run the Hyper-V role and nothing else. With nothing else to run, there's little else that requires RAM on the host.

Not everyone can follow those best practices. Individual needs might require some services or other applications on the host. In this case, you might need to ensure that enough RAM is kept in reserve for host processing. This Host Reserve makes sure that those host-based applications will always have the RAM they need.

You create the host reserve inside the registry. Set the decimal REG_DWORD value for HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Virtualization\RootMemoryReserve to the number of megabytes that should be reserved for the host. The default decimal value is 32, for 32MB, while the maximum value is 1024, 1GB.

—Greg Shields

InstantDoc ID 125975

Q: How can I map a network drive in PowerShell?

A: For use strictly within PowerShell, use the New-PSDrive cmdlet. The resulting drive won't be visible in Explorer. If you want to map a more traditional system wide network drive, use good ol' NET USE, just like you would have done in Cmd.exe.

—Don Jones

InstantDoc ID 125885

Q: Will a PowerShell cmdlet keep going after an error?

A: Some cmdlets can produce non-terminating errors. An example is Get-WmiObject, which will produce one if it can't connect to a computer you've specified. If you specified multiple computers, however, the cmdlet will keep trying the rest of them—it won't just break down and stop. The exact behavior depends on the error itself and the cmdlet you're running.

—Don Jones

InstantDoc ID 125732

Q: Is there an easy way to stop Outlook users from forwarding the messages I send or using Reply to All?

A: Windows Rights Management Services (RMS) has great functionality to strictly control how information is used. In Outlook, it can determine whether users can forward, Reply to All, cut and paste, and more. However, it requires the RMS infrastructure be in place.

Outlook and Exchange have some hidden flags that aren't exposed by the normal UI that can stop Reply to All and forwarding, and Microsoft Research has released an add-in for Outlook 2007 and Outlook 2010 to expose the functions. Download and extract it, run setup.exe, then restart Outlook. Once you restart, you'll have two new buttons in your message composition window for Reply Options. Easy!

—John Savill

InstantDoc ID 126030

Q: With Hyper-V R2 SP1 dynamic memory, how should I set Startup RAM and Maximum RAM?

A: Correctly sizing the amount of RAM you assign to virtual machines (VMs) is a task that nobody does very well. Some of us just assign 4GB and move on. Others assign a minimum amount, adding extra RAM only when users complain.

But there's an argument that determining the correct amount of RAM is best left to the machine itself. Who knows better than that VM how much RAM it needs? Its performance counters always know how much it has and how much it's just about to need.

That's why dynamic memory in Hyper-V R2 SP1 is a pretty slick new feature. With dynamic memory, your Hyper-V host takes care of monitoring memory use and needs on every VM. It does its checks every second, ensuring that immediate VM needs can be quickly taken care of.

This said, using dynamic memory requires you to set values for Startup RAM and Maximum RAM for each VM. These values tell dynamic memory the absolute minimum amount of RAM (Startup RAM) required in order to boot the machine and get the dynamic memory management functionality up and running. The values also identify the maximum amount of RAM that any VM should ever be assigned (Maximum RAM).

You'll obviously then want to keep your Startup RAM at a lowest-possible value. By default, this value is set to 512MB per VM. That value identifies the bare minimum of what is needed to get the machine booted so that dynamic memory can begin identifying its real needs.

Because dynamic memory does all the memory management work for you, each VM's Maximum RAM setting simply needs to be set at a sufficiently high number. The default is 65536MB which gives the manager a supremely large leeway in directing RAM to needy computers.

However, there may come a situation where setting Maximum RAM to a lesser value can be important. For example, if you have a VM workload that will absolutely suck up every bit of RAM that's available for it. Or if you have lots of desktop VMs on a server and you want to prevent a user from loading a monster database.

Setting a much lower limit on Maximum RAM in these cases prevents the situation where one VM still needs the RAM but isn't playing nicely with others in how much it needs. Set this value to a lower quantity when you don't trust the VM not to behave in its memory consumption.

—Greg Shields

InstantDoc ID 125971

"THE CONVERSATION BEGINS HERE"

Microsoft® Exchange CONNECTIONS

WINDOWS CONNECTIONS

SharePoint CONNECTIONS

Microsoft® Visual Studio CONNECTIONS

Microsoft® ASP.net CONNECTIONS

Microsoft® Silverlight CONNECTIONS

SQL Server CONNECTIONS

BONUS: Mobile Apps Track

QUESTIONS ANSWERED • STRATEGY DEFINED • RELATIONSHIPS BUILT

One Place, *One* Time ...

Make **CONNECTIONS** the **CONFERENCE** you bring your whole team to this year!

MARCH 27-30, 2011

ORLANDO, FL

GRANDE LAKES JW MARRIOTT RESORT HOTEL

REGISTER EARLY!

Reserve a room by December 20th *for a minimum of 3 nights* and receive a **$100** Marriott gift certificate, plus a **$100** discount from the regular registration fee.

WinConnections ... Providing the vision + intelligence to keep you and your company competitive in today's market!

Only Microsoft and Industry Experts speak at WinConnections!

Conference Chairs

DON JONES
CONCENTRATED TECHNOLOGY

PAUL ROBICHAUX
TRAINER/AUTHOR

KEVIN LAAHS
HP

KEIRAN MCCORRY
HP

MIKE DANSEGLIO
CONSULTANT

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.WinConnections.com • 800.505.1201 • 203.400.6121 • Register Today!

Microsoft · SharePointPro CONNECTIONS · SQL Server Magazine · Windows IT Pro · TECH Conferences Inc. PENTON MEDIA

COVER STORY

Managing AD in Bulk Using PowerShell

The Active Directory module makes it easy

by Darren Mar-Elia

Regardless of how you manage Active Directory (AD) day-to-day, every so often you'll probably need to perform bulks operations against AD. Whether you're searching for objects that meet particular criteria so you can take some action on them or modifying objects in bulk, the *Active Directory module for Windows PowerShell* is proving to be a worthy tool.

In "PowerShell and Active Directory" (February 2010, InstantDoc ID 103360), I discussed how to load and access the Active Directory module, which is new to Windows Server 2008 R2 and Windows 7. Now I'll show you two examples of how to use the module to automate bulk operations in your environment.

Searching High and Low

One of the most common bulk operations is searching for AD objects that meet particular criteria—say the value of an attribute or the status of a computer account. Let's start exploring the Active Directory module by creating such a scenario.

Suppose I want to find all the disabled computer accounts in three organizational units (OUs) in my AD domain and move them to a temporary holding OU so I can dispose of them later. One way to do this is to create a list of OUs to search, find all the disabled computer accounts in those OUs, and move them as they're found. You might think it would be easier to search the entire domain, but providing a list of OUs reduces the number of objects to search and lets you experiment with providing input to PowerShell commands, which is useful to know how to do.

So, I first need a way to feed the OU names to PowerShell. One approach is to put the names in a comma-separated value (CSV) file, then use PowerShell's Import-CSV cmdlet to read in the names, storing them in a variable. This approach is good if you have a long list of items to feed to a command. However, since I want to search only three OUs, I'll use an array to store their names.

The OU names give PowerShell a place to start its search from, which is called the *search base*. The search base typically takes the form of an LDAP distinguished name (DN)—for example, OU=Marketing, DC=cpandl,DC=com—so that's what I'll store in my array. The following command creates my array of OUs to search:

```
$searchBase =
  "OU=Test,DC=cpandl,DC=com",
  "OU=Sales,DC=cpandl,DC=com",
  "OU=QA,DC=cpandl,DC=com"
```

(Although this command wraps here, you'd enter it all on one line. The same holds true for the other commands that wrap.) This command assigns the string array to the $searchBase variable. If you want to make sure that an array is created, run the command

```
$searchBase[0]
```

PowerShell should return the first element in the array (OU=Test,DC=cpandl,DC=com).

Now I need a command that does the searching. For this particular task, I've chosen the Search-ADAccount cmdlet. The great thing about this cmdlet is that it has all the parameters I need to search for inactive computer accounts, so I don't have to build a complex LDAP filter. For example, to find all inactive computer accounts in a domain, I can simply run

```
Search-ADAccount
  -AccountDisabled -ComputersOnly
```

Figure 1 shows an example of what the results will look like.

Finally, I need a command that moves the disabled computers to a holding OU. To perform the move, I'll use the Move-ADObject cmdlet, which lets you easily move an object or container from one place to another in AD. For example, the following command uses this cmdlet to move a disabled computer account to an OU called BitBucket:

```
Move-ADObject -Identity
  "CN=NT4,CN=Computers,DC=cpandl,
  DC=com"
  -TargetPath
  "OU=BitBucket,DC=cpandl,DC=com"
```

In this case, the -Identity parameter specifies the DN of the object I want to move (a workstation named NT4) and the -TargetPath parameter specifies the DN of the OU I want to move the object to. It's as simple as that.

I now have all the pieces to perform the bulk search-and-move operation: I have the list of OUs, the command to find disabled computer accounts, and the command to move them. There are two ways to put them all together, depending on how comfortable you are with PowerShell.

The simplest approach is to search for inactive computer accounts in a single OU, then use PowerShell's built-in pipelining capability to pipe the results to the command that will move those accounts. This can be done with code such as

```
Search-ADAccount -AccountDisabled
  -SearchBase
  "OU=SDM,DC=cpandl,DC=com" |
  Move-ADObject -TargetPath
  "OU=BitBucket,DC=cpandl,DC=com"
```

In this code, I use the Search-ADAccount cmdlet to search for disabled computer accounts in the SDM OU. I pipe the resulting list of disabled computer accounts to the Move-ADObject cmdlet, which will move those accounts to the BitBucket OU. Note that I didn't need to include the -Identity parameter for Move-ADObject because the pipeline takes care of passing the name of each disabled computer account without me explicitly stating it.

The limitation of this approach is that I have to type this command for each OU on my list. This is where my $searchBase array comes in handy. Using that array, I can write code to iterate through my OU list:

```
foreach ($ou in $searchBase)
  {Search-ADAccount -AccountDisabled
  -ComputersOnly -SearchBase $ou |
  Move-ADObject -TargetPath
  "OU=BitBucket,DC=cpandl,DC=com"}
```

Let's look at what this code is doing. First, I use the ForEach-Object cmdlet (aliased as *foreach*) to iterate through my list of OU names, which is stored in $searchBase. For each OU in $searchBase, I call the Search-ADAccount cmdlet and tell it to look for disabled computer accounts in that OU. I then pipe the output to the Move-Object cmdlet, which moves them. The end result is that all disabled computer accounts in the three OUs are moved to the BitBucket OU.

Making Mass Modifications

The second scenario I want to cover is making mass changes to AD objects. For example, you might need to modify a particular attribute on a large number of objects, based on some other criteria. Let's create a scenario, then look at how to accomplish it.

Suppose I want to find all of the users who are a member of the Marketing Employees group. For each of these users, I want to write the string "FTE" to his or her employeeType attribute.

First, I need to find the best way to determine a user's group memberships. I have a couple of options:

- I could read each user's memberOf attribute to determine whether that

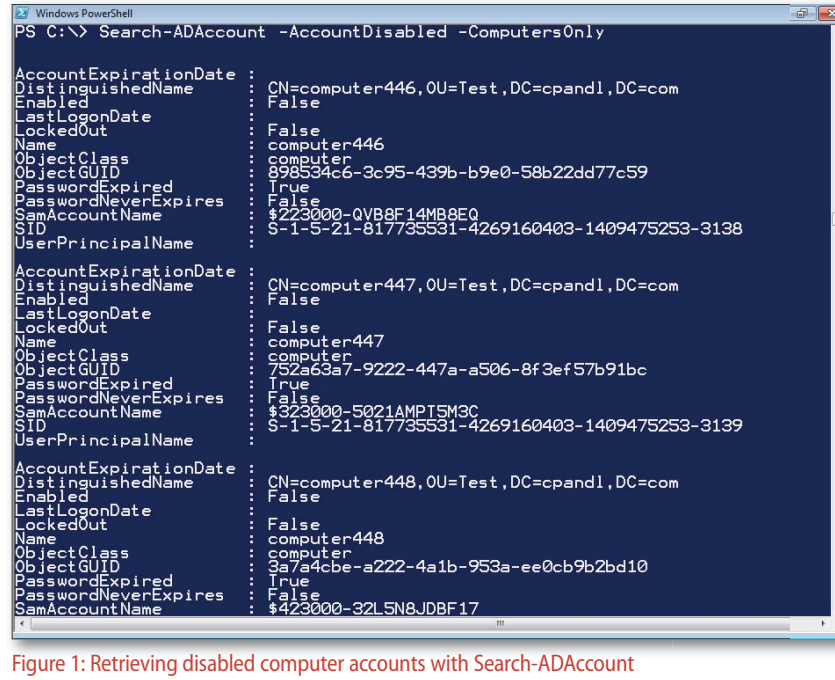

Figure 1: Retrieving disabled computer accounts with Search-ADAccount

user is a direct member of a particular group. That doesn't necessarily give me any indirect memberships (groups that are a member of other groups), but it does get me part of the way there. If I need indirect memberships, I could read a user's tokenGroups attribute, which is a special constructed attribute that represents both direct and indirect group memberships. It exists in Windows Server 2003 AD and later.

- I could search each group, looking at its members attribute to find out that group's direct members. I would also need to search through the membership of each group that's nested within another group. Fortunately, the Active Directory module makes this task easy. Namely, the Get-ADGroupMember cmdlet provides the -Recursive parameter, which will chase down any indirect group members.

The second option is a more scalable solution for my scenario, so that's the approach I'll use.

Next, I need to find the best way to modify attributes on user objects. To do this, I'll use Set-ADUser. This cmdlet lets you modify properties on user accounts. It comes with a set of named parameters that include commonly modified properties. However, you can also modify many other user account properties by using the generic -Add, -Replace, -Clear, and -Remove parameters.

The employeeType attribute isn't one of Set-ADUser's named parameters, so I'll need to use the generic parameters. When using these parameters, it's important to know that they view an attribute that's not set differently than an attribute that already has a value. So, you need to use the -Add parameter when an attribute isn't set and the -Replace parameter when a value already exists.

Now that I have all the pieces in place to make the bulk change to the employeeType attribute based on user group membership, let's put it all together. Once again, I can leverage the PowerShell pipeline, as follows:

```
Get-ADGroupMember
  -Identity "Marketing Employees"
  -Recursive |
  where { $_.employeeType
  -eq $null } |
  Set-ADUser -Add
  @{employeeType = "FTE"}
Get-ADGroupMember
  -Identity "Marketing Employees"
  -Recursive |
  where { $_.employeeType
  -ne $null } |
  Set-ADUser-Replace
  @{employeeType = "FTE"}
```

Let's look at what these two commands are doing. In the first command, I use the Get-ADGroupMember cmdlet with the -Recursive parameter to get all the direct and indirect members of the Marketing Employees group in my domain. I then use the PowerShell pipeline to send the output to the Where-Object cmdlet (aliased as *where*), which checks whether the group member's employeeType attribute is equal to null (i.e., a value isn't set). If it's null, I pass the member to the Set-ADUser cmdlet to populate that attribute. Because employeeType isn't one of Set-ADUser's named parameters and it doesn't contain a value, I use the -Add parameter to populate it with "FTE". The tricky thing about using a generic parameter is that you have to pass in a hashtable that contains the name of the attribute and its value. A hashtable is simply a key-value pair that you can define by delimiting it with the *@{ }* construct, as I've done with *@{employeeType = "FTE"}*.

The second command is basically the same as the first, except that I change the *where* clause to check whether employeeType is not equal to null (i.e., already has a value). If it's not null, I pass the member to the Set-ADUser cmdlet. This time, I use the generic -Replace parameter to change the existing value to "FTE". Once again, I use a hashtable to provide the new value.

Get the Job Done Quickly and Easily

As the two examples show, searching for AD objects is simple with Search-AD Account and moving them is effortless with Move-ADObject. Get-ADGroupMember makes it easy to recursively root out all direct and nondirect members of a group—a task that used to take quite a bit of work in the pre-PowerShell days. And when you need to modify those group members, Set-ADUser quickly gets the job done.

All in all, when it comes to performing bulk operations against AD, the Active Directory module delivers. Its cmdlets provide a powerful built-in mechanism to perform most automation tasks. I encourage you to explore the other cmdlets in the module, which you can enumerate by typing

```
Get-Command -Module ActiveDirectory
```

from the PowerShell console.

InstantDoc ID 125980

Learning Path

WINDOWS IT PRO RESOURCES

To learn more about the Active Directory module for Windows PowerShell, see

"And the Preferred Set of AD Cmdlets is . . ."
windowsitpro.com/go/PreferredADcmdlets

"PowerShell and Active Directory,"
InstantDoc ID 103360

To learn more about how to run and troubleshoot PowerShell scripts, see

"Debugging in Windows PowerShell,"
InstantDoc ID 125694

"Editing and Debugging Scripts with PowerShell 2.0's Integrated Scripting Environment,"
InstantDoc ID 104713

"Error Trapping and Handling in PowerShell,"
InstantDoc ID 125327

"Running PowerShell Scripts Is as Easy as 1-2-3,"
InstantDoc ID 103427

To get your PowerShell questions answered, go to

PowerShell FAQs, www.windowsitpro.com/FAQs/PowerShellFAQs.aspx

Darren Mar-Elia
(dmarelia@windowsitpro.com) is a contributing editor for *Windows IT Pro* and is CTO and founder of SDM Software (www.sdmsoftware.com). He maintains a Group Policy resource website (www.gpoguy.com) and is coauthor of *Microsoft Windows Group Policy Guide* (Microsoft Press).

Smarter technology for a Smarter Planet:

What 27,383 computations per second mean to this energy meter.

They mean this meter will be read 24 times per day, instead of once per month. Giving consumers more visibility into their energy consumption and utility companies a deeper understanding of how energy is being used. eMeter worked with IBM and is using Power Systems™ and IBM application and service management software to enable utilities to manage data from over 20 million smart meters making readings every hour—or more than four times the scale of other utilities industry benchmarks.[1] A smarter business is built on smarter software, systems and services.

Let's build a smarter planet. ibm.com/meter

IBM

A visualization of the data from eMeter's readings for an average home for one year.

1 Based on published benchmark results. Results as of 9/13/10. Sources: IBM press release http://www-03.ibm.com/press/us/en/pressrelease/29315.wss and eMeter press release http://www.emeter.com/2009/emeter-demonstrates-industry%E2%80%99s-most-scalable-smart-grid-management-capability/. IBM, the IBM logo, ibm.com, Power Systems, Smarter Planet and the planet icon are trademarks of International Business Machines Corp, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

Windows Deployment Services in Server 2008 R2

No more installing client OSs by hand

by Rhonda Layfield

PROBLEM:

You need to deploy many clients.

SOLUTION:

Use Windows Deployment Services.

WHAT YOU NEED:

Windows Deployment Services, OS images, an Active Directory domain, DNS, DHCP, and an NTFS partition.

STEPS:

1. Install WDS
2. Configure WDS
3. Add boot images
4. Add install images
5. Add drivers to your images
6. Deploy to clients

Windows Deployment Services (WDS) is Microsoft's replacement for Remote Installation Service. WDS has been around for a while—you could've installed WDS on Windows Server 2003 SP1, and it ships in the box with Windows Server 2008 and later Windows Server OSs. In this article, I'll begin with installing WDS and walk you through configuring WDS, adding images (both .wim and .vhd formats). I'll also cover a topic that makes the top 10 most difficult things to do with Microsoft deployment tools list every time—driver management. I'll finish up by showing you how to deploy an image to both known and unknown clients, and what the difference is.

There are two types of WDS servers: transport and domain-based. WDS transport server was designed for smaller environments that don't have Active Directory (AD) domains. Transport servers require less infrastructure than domain-based servers, but they're more difficult to set up and configure. In this article, I address only domain-based WDS. Domain-based WDS requires an infrastructure that includes an AD domain, DNS, DHCP, and an NTFS partition.

Step 1: Installing WDS

How you install WDS depends on your server's OS. In Server 2003 SP2, you use the Control Panel Add/Remove Programs applet, then access Windows Components, WDS. On a Server 2008 R2 or Server 2008 server, you add it as a role through Server Manager by following these steps:

1. Open Server Manager from the task bar (or the Start menu under Administrative Tools) and click Roles.

2. Click Add Roles to launch the Add Roles Wizard.

3. On the Before You Begin page, click Next.

4. The roles are listed on the Server Roles page. Select the Windows Deployment Services check box and click Next.

5. The Overview of Windows Deployment Services page explains a little about WDS and provides links for configuring and managing WDS. Review the links or just click Next.

6. The Select Role Services page lists two services that need to be installed for a domain-based WDS server. The Deployment Server and Transport Server check boxes should already be selected, so accept the default selections and click Next.

7. On the Confirm Installation Selections page, click Install. You'll see the Installation Progress page, then, when it's finished, the Installation Results page. After it's successfully installed, click Close and you're done—no reboot required. The newly installed WDS snap-in is found in the Start menu under Administrative Tools, Windows Deployment Services.

Before you can deploy images from your WDS server, you need to configure WDS. Open the WDS snap-in, then expand the Servers node and you should see your WDS server's name. If your server's name doesn't appear, right-click the Servers node and choose Add Server. The local computer is selected by default, so clicking OK adds the local server to the WDS snap-in. This is also how you can add multiple WDS servers to a single WDS snap-in so they can be remotely managed from a central location. To add a different WDS server instead of the local computer, right-click the Servers node and choose *Another computer,* then click Browse and navigate to the WDS server you'd like added.

Step 2: Configuring WDS

Now that your server is listed under Servers, notice the yellow yield sign next to your server's Fully Qualified Domain Name (FQDN). Configure WDS by following these steps:

1. Right-click your server name and choose Configure Server to launch the Windows Deployment Services Configuration Wizard.

2. The Before You Begin page lists the requirements of a domain-based WDS server. Click Next.

3. On the Remote Installation Folder Location page, type a path where you'd like to store your images (this must be an NTFS partition). I recommend putting your images on a drive other than your system or boot drives, and you should have at least 20GB of storage space. I always accept the default name of the folder (RemoteInstall) and change the drive letter. Click Next.

4. The page displayed next depends on whether the DHCP service is installed on the same server you're installing WDS. Clients find both DHCP and WDS services by broadcast traffic destined for port UDP 67, which introduces two problems: Broadcast traffic is usually prevented from traversing one subnet to another, and DHCP and WDS installed on the same server can't both listen on port UDP 67. This leads you to three scenarios—WDS and DHCP installed on the same server, WDS and DHCP on different servers but the same subnet, and WDS and DHCP on different servers and different subnets.

- **WDS and DHCP installed on the same server:** WDS and DHCP can't both listen on port UDP 67. Configuring WDS to *not* listen on UDP 67 leaves the port available for DHCP traffic and setting DHCP option 60 (scope or server option) to *PXEClient* resolves this problem. If you're installing WDS on a DHCP server, the DHCP Option 60 page appears during configuration. Select both the *Do not listen on port 67* and *Configure DHCP option 60 to PXEClient* check boxes. When a client receives an offer from that DHCP server, the offer will contain an IP address, subnet mask, and option 60, telling the client that the DHCP server is also a WDS server. If you add DHCP to a WDS server later you can configure this option by right-clicking your server name in the WDS snap-in and choosing Properties—the DHCP tab has similar check boxes.
- **WDS and DHCP on different servers, same subnet as clients:** No additional configuration is required because WDS and DHCP are on different servers. They can both listen on port UDP 67, and the clients are on the same subnet as the DHCP and WDS servers, so broadcast traffic can be heard by all.
- **WDS and DHCP on different servers, different subnets:** When WDS and DHCP are installed on different servers, there's no issue with them both listening on port UDP 67. The problem is that broadcast traffic isn't normally allowed to pass from one subnet to another. There are a couple of ways to resolve this. You can configure IP helpers on your routers and switches to forward the traffic to the appropriate server, or you can configure DHCP options 66 and 67. Configuring DHCP options on an IPv4 network is performed in the DHCP snap-in, found in the Start menu under Administrative Tools, DHCP. From within the DHCP snap-in expand IPv4, then right-click Server Options (or scope options) and choose Configure Options. Scroll down to 066 Boot Server Host Name, select the check box, and type your WDS server's host name in the String value box (to find your server's host name open a command prompt and type *hostname*). Then select the 067 Bootfile Name check box. In the String value box, type the path and name of the Preboot Execution Environment (PXE) file the client will boot. For example, on my WDS server, I accepted the default name of the WDS folder (RemoteInstall), so the string value for option 67 is E:\RemoteInstall\Boot\x86\pxeboot.com.

5. The PXE Server Initial Settings page controls which client computers the server responds to. There are four options. The *Do not respond to any client computers* option is useful if you're not ready for the clients to deploy an image yet (maybe you haven't finished adding your images to the WDS server). The *Respond only to known client computers* option requires that you prestage your computer objects in the Active Directory Users and Computers snap-in. (For detailed steps, see the web-only sidebar "Prestaging a Client" aw www.windowsitpro.com, InstantDoc ID 125868). The *Respond to all client computers (known and unknown)* option is the least secure and, by default,

allows any client that can connect to your network and authenticate to your domain the ability to deploy an image from your WDS server if the permissions on your images are left at the defaults. The last option, *Respond to all client computers (known and unknown) BUT - Require administrator approval for unknown computers,* is my favorite. This option allows known clients (prestaged clients) access to your WDS images and makes you aware of any unknown clients that attempt to access images on your WDS server. I'll explain the difference between known and unknown clients in a later section. I recommend choosing the last option (you also have to select the *Require administrator approval for unknown computers* check box) and clicking Next.

6. You'll see the Task Progress page and then the Operation Complete page. By default, the *Add images to the server now* option is selected. I like to clear this setting and add my images when I'm ready, so clear the checkbox and click Finish. You should now see a green arrow next to your server name.

Your WDS server is now configured. The next step is to add your images. There are two types of images: Boot (Windows Preinstallation Environment, WinPE) and Install (these are the OS images you'd like to deploy).

Step 3: Adding Boot Images

A boot image is what you'll boot your clients with so they can get on the network and be ready to deploy an install image. I recommend using the latest boot image from Microsoft (boot.wim, found on the Windows 7 DVD or .iso in the Sources folder). When you boot using an x86 (32-bit) version of boot.wim, you'll have both x86 and x64 install images listed to choose from. If you boot using the x64 (64-bit) version of boot.wim you'll have only x64 install images to choose from. Follow these steps to add a boot image:

1. Within the WDS snap-in, right-click Boot Images and choose Add Boot Image to launch the Add Image Wizard.
2. On the Image File page, click Browse and navigate to your Windows 7 boot.wim file. Click Next.
3. On the Image Metadata page, accept the default image name and description (or type your own) and click Next.
4. The Summary page shows your selections. If everything is correct, you can click Next.
5. When the Task Progress page appears, click Finish and your new boot image will be listed in the details pane.

Step 4: Adding Install Images

The Install image contains the OS you want to deploy. The supported OSs to deploy are Windows 7, Server 2008 R2, Server 2008, Windows Vista, Windows XP, Server 2003 R2, and Server 2003. WDS supports both .wim and .vhd OS image formats. You can add a Microsoft OS image file or one you've

created. If you create your own image using ImageX, be sure to use the /flags switch so WDS knows which edition of the OS is in the image. The imageX command syntax looks like this:

```
ImageX /Capture <volume to capture>
  <path and name of new .wim>
  "Description" /Flags "Edition"
```

Enter the following command to create an image of the C partition, name it NewImage.wim, give it a description of New Windows 7 Image, and identify the OS as Ultimate edition:

```
ImageX /Capture C: C:\NewImage
  .wim "New Windows 7 Image" /Flags
  "Ultimate"
```

Adding a .wim Install Image to WDS

In this example, I'm deploying a Windows 7 OS image file from the Windows 7 DVD's Sources folder. It's named install.wim. Follow these steps to add your Windows 7 install.wim file to WDS:

1. From within the WDS snap-in, right-click Install Images and choose *Add Install Image* to launch the Add Image Wizard.
2. On the Image Group page, accept the default, *Create an image group named ImageGroup1,* or type a new image group name. I'm deploying an x86 Windows 7 image, so I'll create an image group named Windows 7 x86, then click Next.

Image groups reduce the storage space needed for your images. I'm adding a Windows 7 32-bit OS image to the image group named Windows 7 x86. When a second 32-bit Windows 7 image is added to the same image group, a feature called single instancing looks at every file within the second image that's being added to the image group. If the file already exists (because it was in the first image I added), a pointer is created to use the existing file and not store the same file again. There are a lot of core OS files that only need to be stored once within an image group, so you can save substantial space. The more images you add, the greater the savings. Because the 64-bit core OS files are completely different from the 32-bit architecture, I recommend creating one image group for 32-bit images and another for 64-bit images. Also note that single instancing only works with .wim image files, not Virtual Hard Disk (.vhd) files.

3. On the Image File page, click Browse, navigate to your Windows 7 DVD's Sources folder, and double-click install.wim. Click Next.
4. The Summary page shows the images that are contained in install.wim file. I'm using the retail DVD's install.wim file, which contains Windows 7 Starter, Home Basic, Home Premium, Professional, and Ultimate editions. Click Next.
5. The Task Progress page appears. Because install images are normally much larger than boot images, this will take longer to add. When the install image has been added, click Finish. The new install image appears in the details pane.

Adding a .vhd Image to WDS

Adding a .vhd image is different than adding a .wim file in that it can't be done within the WDS snap-in. You must use the WDSUtil command-line utility to add a .vhd file. For example, if I have a .vhd file named Win7.vhd stored in the C:\Images folder and want to add it to WDS's image group named MyVHDs, I'll follow these steps:

1. From within the WDS snap-in, create the image group MyVHDs by right-clicking the Install Images node and choosing *Add Image Group*.
2. In the Add Image Group box, type *MyVHDs* (or any name you want) and click OK. The new image group appears under the Install Images node. Close the WDS snap-in.
3. Open an elevated command prompt on the WDS server by right-clicking the Command Prompt and choosing *Run as administrator*.
4. Using WDS's command-line utility, WDSUtil, add a .vhd install image with the following commands:

```
WDSUTIL /Add-Image
  /ImageFile:"PathAndNameOfVHD"
  /ImageType:Install /ImageGroup:
  ImageGroupName
```

The command for adding C:\Images\Win7.vhd is:

```
WDSUTIL /Add-Image /ImageFile:"C:\
  Images\Win7.vhd" /ImageType:Install
  /ImageGroup:MyVHDs
```

Step 5: Adding Drivers to Images

Drivers in WDS are by far the most complicated component. They're not well documented and are incredibly frustrating to figure out. I hope these step-by-step instructions ease the WDS driver pain for you.

WDS uses dynamic driver provisioning (DDP) to control which drivers are deployed to your clients. You could rely on your clients to perform Plug and Play (PnP) during deployment to determine which drivers are needed. Or you can control driver deployment by organizing your drivers into driver groups.

A driver group is a collection of driver packages (drivers packaged in .exe, .msi, or .cab formats). You can add filters to the driver group identifying which drivers from the group are available to specific client computers based on hardware (such as manufacturer, BIOS vendor, or version) or attributes of the OS install image such as version (Windows 7 or Server 2008) or edition (Ultimate or Enterprise).

Normally the only drivers that will be deployed are drivers associated with hardware that is currently attached and powered on to your client machines. With DDP, you can install drivers for hardware that isn't yet connected to your clients. It's a common practice to combine driver group filtering and PnP.

The supported driver packages are available only for Windows 7, Server 2008 R2, Server 2008, and Vista SP1. WDS can't read driver packages themselves—they must be extracted first. Drivers packaged as .msi or .exe can sometimes be extracted using WinZip or WinRAR. With some .exe files, you can change the extension from .exe to .zip, unzip the file, and use Windows Explorer to copy the files from the compressed folder. Drivers packaged as .cab can be extracted using the Expand utility. For example, a driver package named videodrv.cab stored in the C:\drivers folder can be extracted to C:\expanddrv (which you created) by opening an elevated command prompt and typing

```
expand -F:* c:\drivers\videodrv.cab
  c:\expanddrv
```

Note that the -F in the expand command is case sensitive.

Adding drivers to WDS is performed in two major steps: first you must add the driver packages to the WDS snap-in which also associates the driver package to an install image group (all install images within the image group have access to the driver group), then you can add the drivers to a boot image. To add driver packages to the WDS snap-in right-click the Drivers node, click Add Driver Package to launch the Add Driver Package Wizard:

Adding Driver Packages to WDS

1. On the Driver Package Location page, choose to either add a driver package from an .inf file or a folder containing your expanded driver packages. Then click the Browse button and navigate to the .inf file or folder and click Next.
2. The drivers you added in step 1 are listed on the Available Driver Packages page. All drivers are selected by default. If you'd like to de-select a driver, remove the checkmark next to the driver. After selecting your drivers click Next.
3. The Summary page lists the drivers to be added. If everything looks good, click Next.
4. The Task Progress page displays each driver as it's added. When it's done, click Next.
5. There are three choices for adding drivers to driver groups on the Driver Groups page: add drivers to an existing group (select your driver group from the drop down list), create a new driver group or don't put the driver in a driver group now. I created a driver group named VMware and clicked Next.
6. On the Task Complete page accept the default setting to modify the filters for this group. Click Finish to open the properties of the new VMware driver group.
7. To add a filter, click the Add button. Select the type of filter (I chose Manufacturer, as Figure 1 shows), select an Operator (I chose *Equal to*), type the name of the manufacturer (such as VMware, Inc.), and click OK.
8. The new driver group appears under the Drivers node and the drivers are displayed in the details pane. The information displayed includes Package Name, File Name, and Class GUID. You might need this information when adding your drivers to a boot image later.

Adding Drivers to a Boot Image

Now that your images have been added to the WDS snap-in, you can add drivers to your boot images for your client machines' hardware. What type of drivers do you need to add? Possibly mass storage or network drivers. Follow these steps within the WDS snap-in:

1. Expand the Servers node, then expand your server and click on the Boot Images node. In the details pane, right-click the image and choose *Add Driver Packages to Image* to launch the Add Driver Packages to Image Wizard.
2. The Before You Begin page warns you to first back up your boot image by right-clicking the image and choosing Export Image. Exporting the image won't remove it from the WDS snap-in, it just creates a copy. If you don't have a backup copy of the boot image, click Cancel and export the image first, then launch the wizard again. If you have a backup copy of your boot image, click Next.
3. Two default filters will be displayed on the Select Driver Packages page. The filters will depend on which boot image you launched the Add Driver Packages to Image Wizard from and its architecture. I right-clicked a x86 boot.wim file, so my filters were *Package Architecture equal to x86* and *Package Class equal to "Net," "System," "DiskDrive," "hdc," "SCSIAdapter."* You can add more filters, or if these are the only filters you need, click the Search for Packages button. All drivers that meet the filter criteria will be listed in the Search results box. By default all drivers that meet the criteria will be selected, accept the default drivers or deselect the ones you don't want and click Next.
4. The drivers you chose will be listed on the Selected Driver Packages page; click Next.
5. The Task Progress page shows that the boot.wim file is mounted, the drivers are added (each driver is listed as it is added), changes to the boot.wim are saved, and it's unmounted. On the Operation Complete page, click Finish.

Viewing the Drivers in a Boot Image

I wish there were an easy way to view the drivers that have been added to the boot image. To view the drivers in a boot image follow these steps:

1. Export the boot image by right-clicking the image and choosing Export Image. Give the exported image a name and path and click Save. I named mine WDSBoot.wim and saved it in F:\wims .
2. Open an elevated command prompt and mount the image to an existing empty folder. Use the Deployment Image Servicing and Management tool (DISM) to mount an image. The syntax looks like this:

```
Dism /Mount-Wim /WimFile:<path and
  name of boot image.wim>
  /index:<index number>
  /MountDir:<path to an existing empty
  folder>
```

For example, to mount my WDSBoot .wim file to the c:\mount folder, type:

```
Dism /Mount-Wim /WimFile:f:\wims\
  wdsboot.wim /index:2 /MountDir:c:\
  mount
```

Mounting boot images doesn't take long—they're small compared with install images.

About the /index:2 switch in the previous command: Microsoft .wim image files can contain multiple images, and by default the boot.wim file I added from the Windows 7 DVD contains two images. Drivers are added to the second image of the boot.wim file, so thats the image you need to mount.

3. To get a list of currently installed drivers from the mounted boot image, type

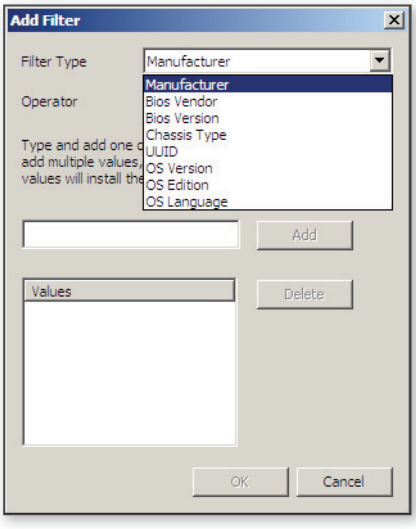

Figure 1: Selecting a filter type

```
Dism /image:c:\mount /Get-Drivers
```

If you have a lot of drivers and don't want the screen scrolling all over the place, you can also get the driver listing in a table format by typing

```
Dism /image:c:\mount /Get-Drivers
  /format:table
```

A great way to quickly document which drivers are currently installed is to pipe this output to a text file. To pipe the information to a text file named drivers.txt and stored in the C:\ drive, type

```
Dism /image:c:\mount /Get-Drivers
  /format:table >c:\drivers.txt
```

4. Next you need to unmount the image. If all you did was view the currently injected drivers, there's no need to save your changes. But if you used DISM at this point to inject other drivers or modify any other part of the boot.wim, you need to save your changes when you unmount. To unmount the image with no new changes, use

```
DISM /Unmount-Wim /MountDir:C:\mount
  /discard
```

To unmount the image if you made changes:

```
DISM /Unmount-Wim /MountDir:C:\mount
  /commit
```

5. If you made changes to the boot.wim file using DISM, need to add the boot image back to the WDS snap-in in one of two ways. From within the WDS snap-in, either right-click the boot image you exported and choose Replace Image, then select the image you previously exported, or add the boot image as a new image, as I did earlier.

Step 6: Deploying to Clients

Deploying your first WDS client is performed by PXE booting, also known as a network boot. When the client boots, it gets an IP address (and any options) from a DHCP server. Next, the client finds a WDS server (also called a Network Boot Server) to boot a Network Boot Program from. If there's more than one boot image on the WDS server, a list of boot images is presented to choose from. Lastly, a list of OS images are presented to choose from. In this section I'll show you each step in more detail and explain the difference between a known and unknown client deployment scenario.

Known Clients

1. To perform a network (PXE) boot from the client, boot a client then, when prompted, press F12 for a network service boot.

2. If only one boot image has been added to your WDS server (as in this example), that boot image is automatically booted. If you have more than one boot image, a menu will allow you to choose which image you'd like to boot.

3. The client boots to the Windows Deployment Services Wizard and prompts you for the installation language. Accept the defaults for U.S. English or choose another language and click Next.

4. You need to authenticate to the WDS server in the authentication dialog box (the user account you authenticate with needs to have read and execute permissions on the image you want to deploy). Be sure to use the *DomainName\UserAccount* format. Click OK.

5. When the list of install images is displayed, select the install image you want deployed and click Next.

6. On the *Where do you want to install Windows* page, click *Drive options (advanced)*, then New.

7. Accept the default size or change it to what you want and click Apply. A message letting you know that Windows might create additional partitions for system files appears. Click OK. By default, deploying Windows 7 creates two partitions. One 100MB partition is created in case you want to use BitLocker (this partition will be hidden and receives no drive letter). The other partition uses the rest of the hard drive and will become the C: partition. Click Next.

8. The *Waiting for server* page appears, followed by the Installing Windows page. When it's done, the client reboots and displays the Set Up Windows page. Choose your country or region, time and currency, and Keyboard layout and click Next.

9. Enter the username you want to create locally on the client and click Next.

10. On the *Set a password for your account* page, type a password (twice) and a password hint and click Next.

11. On the *Type your Windows product key* page, you can enter your product key. If you leave this field blank, you'll have 30 days to provide the product key. (I leave the product key blank on test machines.) The *Automatically activate Windows when I'm online* check box is selected by default. If you choose not to enter your product key, clear this check box and click Next.

12. The *Please read the license terms page* requires you to select the *I accept the license terms* check box then click Next.

13. On the *Help protect your computer and improve Windows automatically* page, there are three selections: *Use recommended settings* (downloads and installs all OS patches, drivers from Windows Update and Windows Defender updates, IE phishing filters, and other updates when they're available), *Install important updates only* (downloads and installs critical updates only), and *Ask me later* (turns automatic updates off). Choose the setting that's right for you and click Next.

14. Set your time zone, date, and time on the *Review your time and date settings* page and click Next.

15. Setting your computer's current location determines how secure your new client will be. If you're on a trusted home network (that has a firewall between the client and the Internet), choose *Home network*. If the client is on a trusted network at work (again, where there's a firewall present) choose *Work network*. If you're at

Figure 2: Screen displayed when performing a network boot

a coffee shop or other public location with no known firewall between the client and the Internet, choose *Public network*.

16. You'll see *Preparing your computer*, then the Windows 7 desktop.

Unknown Clients

The only difference between known and unknown clients is only Step 1. At the beginning of this article in the Configure WDS section (step 5), I recommended that you set the PXE Server Initial Settings to *Respond to all client computers (known and unknown)* and also select *Require administrator approval for unknown computers etc.*

When unknown computers (not pre-staged in the Active Directory Users and Computers snap-in) attempt to perform a network boot (PXE) they receive the screen that Figure 2 shows. Notice at the bottom of this screen there's a message, Pending Request ID: 4, a space to provide a message from the administrator, and a contacting server IP address (this is the IP address of the WDS server the client is attempting to contact). The client will continue to contact the WDS server every few seconds to see if the request has been approved or rejected. Every time the client checks with the WDS server, another dot appears after the WDS server's IP address.

To find pending devices, follow these steps on the WDS server:

1. Open the WDS snap-in and expand Servers and your server name. Then click the Pending Devices node; in the details pane you'll see a pending device with an associated number. In our case it would be number 4.

2. Right-click the pending device Request ID 4. From the menu there are three options: Approve (approves the client; the client will continue with the deployment just as a known client would), Name and Approve (also approves the request, but allows the administrator to name the computer object that's created in ADUC for the new client), and Reject (reboots the client).

From this point forward, all steps for an unknown client are the same as for known client computers. Follow Steps 2-16 to complete your client deployment.

I know there's a lot of information in this article, with installation, configuration, adding basic boot, and install images (along with their drivers) to the client deployment process for both known and unknown clients. I haven't even told you about multicasting transmissions and the advanced configuration options. Look for future articles that cover these and many more WDS topics.

InstantDoc ID 125867

Rhonda Layfield
(rhonda@deploymentdr.com), a Setup and Deployment MVP, offers lecture-based and hands-on deployment classes. For Rhonda's deployment class schedule, please see "Where is Rhonda" at www.DeploymentDr.com.

Liberating Desktop Virtualization

Quest® vWorkspace. Master virtual desktop and application delivery through a single user access point and management console. vWorkspace blends Terminal Server/Remote Desktop Session Host, VDI, Blade/Physical PCs, and Application Virtualization into one solution. With the added freedom to choose from multiple virtualization platforms, vWorkspace delivers simplicity through consolidation. Get more with less complexity, resources and cost.

Experience how Quest liberates Desktop Virtualization Management. Read the white paper "Concept Becomes Reality With Quest" or watch the video at **quest.com/Liberating**.

QUEST SOFTWARE®
Simplicity At Work™

Windows Server 2008 R2 SP1 and Hyper-V: More than a Bunch of Fixes

With its improvements to virtualization, this service pack adds a lot

by John Savill

When you think *service pack*, you probably think of minor updates and fixes. You think of phrases such as "increased reliability" and "performance gains," not "industry-leading new features." These perceptions would be accurate if I were talking about Windows 7 SP1. It has updates to resolve problems and improve performance, and some minor feature updates related to third-party federation integration, HDMI audio performance, and XPS document rendering, but it's nothing to justify an article. When you look at SP1 for Windows Server 2008 R2, however, it's a different ballgame. To say it adds industry-leading new features is no exaggeration. SP1 is a game changer for virtualization, particularly Virtual Desktop Infrastructure (VDI).

SP1 for Server 2008 R2 makes some minor updates outside of Hyper-V, but in this article I'll be looking at the Hyper-V changes. Even before SP1, the hypervisors from different vendors had nearly reached parity, with little difference between them in terms of performance and functionality. You made virtualization choices based on price, management, and integration with the rest of an organization's infrastructure. Hyper-V, however, had one weakness: It lacked memory overcommitment. With SP1, Microsoft has addressed this deficiency by avoiding memory overcommitment like the plague. Confused? Read on.

Hyper-V 2008 R2 SP1 introduced dynamic memory, which allows you to define an initial amount of RAM for a VM and the maximum amount of RAM it can be allocated. Hyper-V intelligently allocates memory to VMs over their initial amounts based on need and on the amount of physical RAM that's available. This is different from memory overcommitment, where you start each VM with the maximum amount of memory possible regardless of if or how it's being used and hope you don't run out of resources. With dynamic memory, VMs are allocated additional memory if it's available, and memory can be reallocated from other VMs that need it less.

Figure 1 shows the dialog box for configuring memory for a Server 2008 R2 SP1 hosted VM. Notice that you can still use the old Static configuration, where you assign a VM a set amount of memory that is all allocated when the VM is turned on and can't be increased. The more interesting option is the Dynamic selection as you see selected in the figure. Startup RAM is the memory allocated to the VM when it is initially turned on and Maximum RAM is the size the memory for the VM can grow to based on the VM's needs and physical memory availability. The default value for the Maximum RAM is 64GB (the maximum supported by a VM in Hyper-V). I suggest you set more realistic values, both for planning and to protect you from some rogue process in one VM that allocates as much RAM as it can. In my example, I've set the maximum to 2GB. Based on my expected workloads, this is a reasonable amount for the VM.

You can also see two sliders in the dialog box. The first is a percentage of memory to keep as a buffer. You set a buffer because you don't want an OS to totally run out of memory before Hyper-V starts giving it more. The process of adding additional RAM can take a few seconds, and during those few

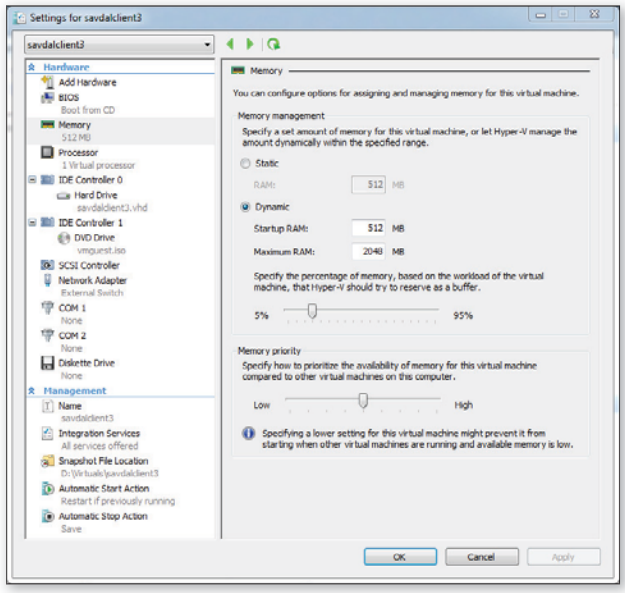

Figure 1: Configuring dynamic memory for a VM

seconds the performance of the VM could be crippled—the guest OS would start to move memory pages to its page file to handle the lack of RAM. To avoid this memory starvation, you set a desired percentage of memory to always be available in the VM (20 percent by default), and when the VM has less than that percentage available, more memory is added to the VM to bring it back to the desired figure (assuming RAM is available in the host). If 20 percent is too little or too much based on the VM's needs, you can change this setting using the slider, but 20 percent is generally a good value for most configurations.

The other slider is to set the priority of memory allocation for when there isn't enough physical RAM available to provide the desired amounts to all the VMs. Just like with CPU allocation, VMs with higher memory priority will receive additional memory before VMs with lower priority.

I've said that dynamic memory intelligently allocates additional memory to a VM, and I use the phrase available memory, not free memory. This is key because available memory and free memory are very different. Windows Vista and later OSs use all the memory they can for caching, helping improve performance by preloading programs into memory. The memory used for this caching can be used by applications whenever it's needed, so the cache is largely still available. Looking at free memory is fairly meaningless—you need to consider the available memory, which includes most of the memory being used for cache. This is exactly what dynamic memory does.

An update to the Hyper-V integration services provides a new dynamic memory Virtual Server Client (VSC) in guest OSs. The VSC communicates with its corresponding Virtual Service Provider (VSP) in the parent partition to report its use of memory, specifically its amount of available memory. Based on the amount of available memory in the guest, the desired memory buffer configured for the VM, and the amount of physical RAM available in the host, additional memory may be allocated to the guest.

This type of intelligent memory allocation is only possible because of the guest OS insight provided by the dynamic memory VSC. It wouldn't be possible if the hypervisor just looked at what memory is being used by a VM from the outside—Hyper-V wouldn't be able to tell if the memory was being used by an application or just for disposable purposes like pre-caching. (See the web sidebar "Virtualization RAM Technology," InstantDoc ID 126048, for more about the technologies involved.) Figure 2 shows the memory that a VM has allocated and its percentage of available memory. You should see that the available memory is pretty close to the target memory buffer percentage you configure.

Some versions of Windows have long had the ability to hot-add memory to the OS, but dynamic memory doesn't use this capability. Hot-addition of memory was the rare instance where you'd add an entire stick of memory to your hardware. That's very different from frequently adding memory in small amounts, and it was found that memory hot addition wasn't the right solution. Instead, the integration services for Hyper-V that run inside guest OSs were enhanced with a new kernel-level memory enlightenment that communicates with the parent. When it's told that additional memory has been allocated, the integration services present it to the guest OS, which continues working with its increased amount of memory. In my tests, this method has worked great.

When a VM doesn't need some memory any more, or another VM needs it more, dynamic memory uses a balloon driver to reclaim memory. A balloon driver is a kernel mode device driver, so when it asks for memory, the OS has to fulfill the request. Hyper-V tells integration services to grow the balloon driver to a certain size. The balloon driver demands the memory from the guest OS and leaves it free for Hyper-V to reallocate. The guest OS can intelligently decide where that memory will come from, including moving data with the least need to be in memory to the local page file. If the guest needs more memory later (and it's available), the balloon driver can deflate, returning memory for the guest.

When to be Dynamic

Dynamic memory isn't the right solution for every virtual workload, but it's a benefit to most. You might typically assign some services, such as domain controllers and file servers, 4GB of memory, but when you start paying attention, you could be surprised how little memory they actually use. VDI environments are another great fit for dynamic memory because end-user machines occasionally need large amounts of memory for intensive tasks but can usually get by with a lot less. So when isn't it a good fit? Consider services that are very intelligent about memory and allocate memory when they're started, or services that will always consume as

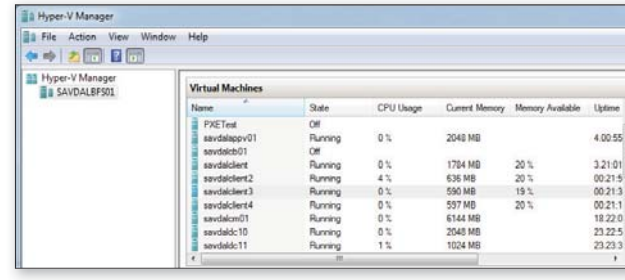

Figure 2: Memory configurations in Hyper-V Manager

much memory as available. (Both SQL Server and Exchange Mailbox servers fit into both of these categories.) If you try to use dynamic memory with these services, they'll just absorb all the memory you throw at them, unless you use service level configurations to limit the amount of RAM that can be used. Generally, a static memory configuration will probably be better than a dynamic one for these memory-absorbent services.

Dynamic memory doesn't mean we don't need to plan. You still need to figure out the normal memory usage patterns of your virtual environments and allocate resources accordingly. Dynamic memory means you don't have to allocate the peak memory used for the entire time a VM is running—you can allow a VM to have more memory when it's needed but use less when it's not. This lets you fit more VMs in your memory, saving you money and management effort.

RemoteFX

RemoteFX is made up of three capabilities, all based around Microsoft's acquisition of Calista Technologies. Calista focused on improving the experience of presentation virtualization technologies such as Remote Desktop. One of the capabilities I'll discuss here is available for Remote Desktop Session Host (formally known as Terminal Services), but the two main capabilities are only available for Windows 7 VDI environments.

RemoteFX vGPU is aimed at providing consistent graphical fidelity for end users, no matter the capabilities of their endpoint device. Without vGPU, users connecting to Windows 7 VDI sessions from Windows 7 clients can enable desktop composition in the remote connection settings and get the full Aero Glass experience. Using multimedia redirection, certain types of supported media, such as WMV files, are sent raw to the client device and rendered locally, giving very smooth multimedia playback. Windows 7 clients get a good experience thanks to the local Windows 7 OS and fairly powerful local hardware. But when connecting from a basic thin client or legacy OS, users get a very different experience—basic graphics only and almost no rich media capabilities. vGPU is about equalizing the experience and letting all users have the same Windows 7 VDI experience.

vGPU works by using a new virtual GPU that is presented to the guest OS through an updated VMBus virtualization driver that is part of the SP1 integration services update. This vGPU uses a GPU in the host to perform the actual graphical computations. Note that you don't need one physical GPU per vGPU in a guest—one physical GPU can support multiple vGPUs in the same way that one physical processor core can be used by multiple virtual CPUs.

You have two options for the GPU requirements in the host. One option is to use graphics cards in the hosts that must support DirectX 9.0c and 10 with at least 256MB of RAM (although you'll want much more than this for any sizable implementation), PCI-Express connections (x16 ideally), and the processors in the host must support Second Level Address Translation (SLAT). Note that in production environments you'd use specialized high-end graphics cards and not consumer-type GPUs, although in a lab environment a more basic GPU would be fine. (I use a GTX 275 with 1GB of RAM in my lab.) Note that these GPUs can be external in the form of an appliance. Your other option is to use a hardware device known as an Application Specific Integrated Circuit (ASIC), which are used to offload GPU and CPU from the host.

With this vGPU, the guest OS has advanced graphical capabilities, including DirectX 3D 9.0c (which is used by productivity applications such as PowerPoint 2010), Silverlight, Adobe Flash, Aero, and many others. All of the rendering and commands associated with the graphics on the VM vGPU are sent via the VMBus to the host OS Render Capture Compress component and replayed on the host GPU using off-screen memory. All rendering is done host side. The capture component then scans for changes in the off-screen memory display rendering, compresses the changes, and sends the changes to the client for display. This RemoteFX graphical content is sent over a new RemoteFX graphical Remote Desktop Protocol (RPD) virtual channel, which is why the client must support RDP 7.1 and be RemoteFX-enabled. Notice that this setup doesn't use the endpoint's rendering capabilities at all, so you get the same experience no matter the capabilities of the local client (assuming it's a RemoteFX-compatible client). If the client isn't RemoteFX-capable or the client isn't connecting via LAN (a LAN connection is a requirement for RemoteFX), the client will default to the standard RDP experience and RemoteFX vGPU won't be used.

The good news is many clients will be able to take advantage of RemoteFX, including standard Windows clients, traditional thin clients, ultra-light thin clients with RemoteFX ASIC components for the decoding and decompression, and even monitors with the RemoteFX ASIC built-in. All of these clients will get exactly the same graphical experience, because all the rendering is performed host side and only screen updates are sent over the network. Remember that the target must be running Windows 7 SP1 on Hyper-V 2008 R2 SP1 with RemoteFX enabled—you can't use the RemoteFX vGPU and associated advanced graphical capabilities on a Remote Desktop Session Host (RDSH).

While I'm talking about the graphical side of RemoteFX, I must mention a component of RemoteFX that isn't talked about that often. RemoteFX has a new set of intelligent codecs for encoding and decoding of the graphical data sent over RDP, compressing the information. These codecs are key to the entire RemoteFX, giving a better experience and using less bandwidth. If you have an ASIC in the RDSH, the codec operations will still benefit by offloading encoding and decoding to the ASIC without using GPU capabilities.

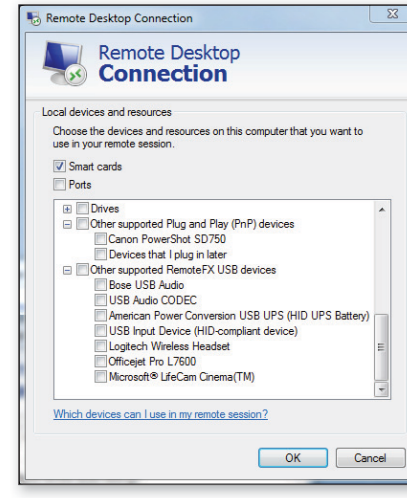

Figure 3: Configuring USB redirection

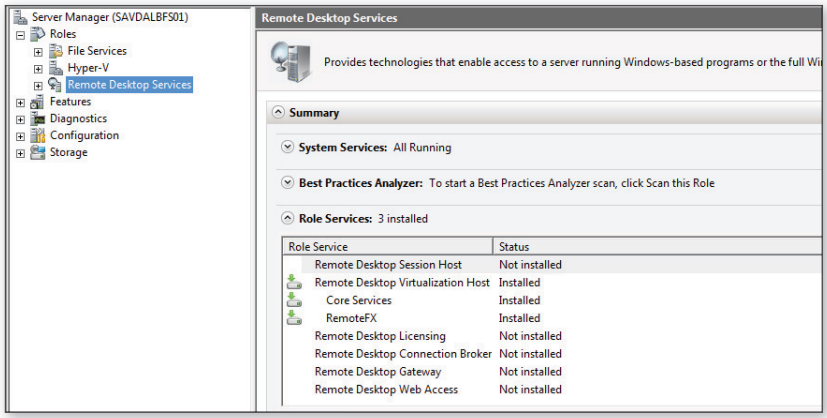

Figure 4: RDS components needed to enable RemoteFX

RemoteFX USB Redirection

The other major feature of RemoteFX is USB redirection. Traditional RDP has several high-level device redirection capabilities, such as for input direction, smart card redirection, port redirection, bi-directional audio, picture transfer protocol, printers (using EasyPrint for driverless printing), and drive redirection. This might sound like a lot but in reality, there are still a huge number of devices that aren't supported by normal RDP redirection.

RemoteFX addresses this by redirecting a device at the USB protocol level. This means far more devices can be redirected using the RemoteFX USB redirection capability, and you get support for devices such as scanners, all-in-one printers, PDAs, mobile phones, multimedia headsets, webcams, and biometric devices. You can see an example of this improved USB redirection support in Figure 3. With normal RDP redirection, I see one device that can be redirected. With RemoteFX USB redirection, I see many more.

The addition of this feature doesn't mean the high-level redirection of normal RDP is no longer needed. With RemoteFX USB redirection, you're redirecting at the USB level, so the driver for the device is redirected and must be installed in the remote OS. The device is no longer available on the local client, which wouldn't be good for keyboards and mice. With normal RDP redirection, the driver is only needed on the local client and that redirected device can be used on multiple sessions, whereas RemoteFX USB redirection makes it accessible to only one session at a time.

In the Server 2008 R2 SP1 release of RemoteFX, USB redirection and vGPU are tied together. If you can't enable vGPU on a VM because your processor lacks SLAT support or your server doesn't have a FPU, you won't be able to use RemoteFX USB redirection. This also means the RemoteFX USB redirection isn't available for RDSH-based sessions. You can expect enhancements in future versions, including a potential decoupling of these capabilities.

Because you're using an RDS component, you need RDS CALs for clients that are taking advantage of RemoteFX. Because it's tied to VDI configurations, however, most companies that could use RemoteFX already have the RDS CALs as part of a suite to enable use of other RDS components commonly used in VDI deployments.

Enabling and Using RemoteFX

For Hyper-V based VMs to use RemoteFX, you have to make several configuration changes to the server, and you should be sure the server meets the RemoteFX requirements, such as having a GPU and SLAT support on the processor. You must enable the Remote Desktop Virtualization Host role service that's part of the Remote Desktop Services role and ensure the Core Services and RemoteFX child component of the Remote Desktop Virtualization Host are installed, as Figure 4 shows.

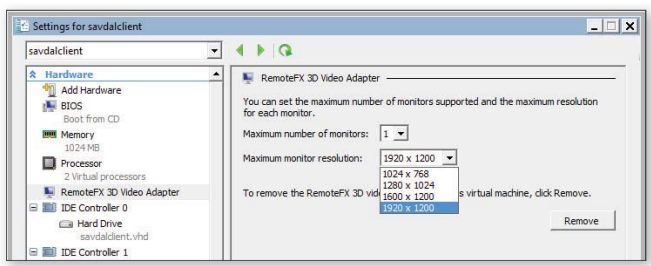

Figure 5: Configuring RemoteFX

Once the component is installed, you'll see a new type of hardware available to add to a VM, the RemoteFX 3D Video Adapter. Add this hardware to the VM and then configure it for the VM, as Figure 5 shows. Ensure the guest is running Windows 7 SP1 and then make your connection from an RDP 7.1, RemoteFX-enabled remote client.

To allow RemoteFX USB redirection, you need to make a change for your clients, a change that's commonly made using Group Policy. Go to the policy Computer Configuration, Administrative Templates, Windows Components, Remote Desktop Services, Remote Desktop Connection Client, RemoteFX USB Device Redirection and set *Allow RDP redirection of other supported RemoteFX USB devices from this computer* to Enabled. Set who has RemoteFX USB redirection rights then click OK and close the policy editor.

Remember that RemoteFX is only for LAN connections today, so check the Experience tab of the Remote Desktop Connection Client. If you choose anything other than LAN, RemoteFX will be disabled and you'll get a normal RDP experience.

SP1 is key in getting the most from your Hyper-V based virtualization infrastructure. SP1 not only catches Hyper-V up with the competition in terms of memory density, but also provides more intelligent handling of memory for virtual environments. With RemoteFX, Hyper-V based VDI environments offer an unmatched experience on any form-factor end-point that supports RemoteFX. These features definitely make SP1 something that you should try out in your environment.

InstantDoc ID 126048

John Savill
(john@savilltech.com) is a Windows technical specialist, an 11-time MVP, and an MCITP: Enterprise Administrator for Windows Server 2008. He's a contributing editor for *Windows IT Pro*, and his latest book is *The Complete Guide to Windows Server 2008* (Addison-Wesley).

Prime Your Mind

with Resources from Left-Brain.com

Left-Brain.com is the online superstore stocked with educational, training, and career-development materials focused on meeting the needs of IT professionals like you.

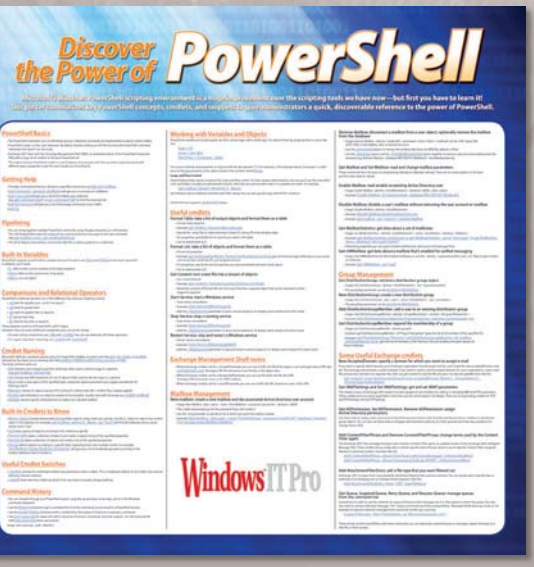

Featured Product:

Windows PowerShell Poster
Discover the Power of PowerShell

Microsoft's Windows PowerShell scripting environment is a huge improvement over other scripting tools, and we can help you learn it! Our new PowerShell poster summarizes key PowerShell concepts, cmdlets, and snippets for group management, Exchange, and other admin tasks.

Topics covered are PowerShell basics, pipelining, built-in variables, mailbox management, command history, and much more!

Only $14.95*!

Order your poster and discover other great PowerShell resources now at Left-Brain.com

left-brain

*Plus shipping and applicable tax.

www.left-brain.com

Windows IT Pro

FEATURE

Inside the Ops Manager Management Pack

A closer look at management pack design, function, and tuning

by Pete Zerger

System Center Operations Manager 2007 R2 is touted as the best-of-breed platform for application monitoring on Windows OSs. The "secret sauce" that separates Ops Manager from the competition is the management pack. Management packs provide the instructions for Ops Manager to discover and monitor specific applications. In June 2008, Microsoft declared that the company would deliver a management pack for every server application the company released, thus ensuring improved manageability across all Microsoft server applications. In this article, I explore the life cycle of a management pack, including design, function, and tuning.

Definition

A management pack is an XML file that contains all the elements necessary to discover and monitor an application with Ops Manager 2007. The management pack is imported into an Ops Manager 2007 management group and distributed to appropriate agents that use its contents to perform monitoring activities for a particular application.

Management packs contain product information, included by the management pack author, that provides common causes and resolutions for alerts that are raised during monitoring. This product information transfers the application experts' (i.e., the Microsoft application architects') knowledge to the operators responsible for responding to alert conditions. Management packs for Microsoft applications are designed primarily by the product teams themselves—the architects who arguably know more about the application than anyone—which gives Ops Manager an advantage over other enterprise monitoring platforms.

Management Pack Components

Management packs contain several components that help discover, describe, and monitor applications. I discuss the following components and related concepts later in the article.

- Attributes
- Classes
- Groups
- Health state
- Monitors (unit monitors, aggregate rollup monitors, dependency rollup monitors)
- Object discoveries
- Overrides
- Rules
- Tasks
- Views
- Reports

The basic components that comprise a management pack include the attributes, classes, groups, health state, and monitors. These elements (listed here in alphabetical order) are viewable in the Authoring space of the Operations console.

- Attributes: Attributes contain information that further defines an object type in Ops Manager. The Attributes node in Ops Manager's Authoring space displays a list of attributes for each object type in the management group, such as Windows Computer, Windows 2003 Operating System, SQL Server Database, or IIS Web Application.
- Classes: A class represents a type of object. Classes are defined in management packs and aren't limited to computers and groups. Classes can be defined to represent any component of an application or device, such as a SQL Server instance or database. Classes are sometimes called *object types* or *targets*.
- Groups: Groups are used to create collections of objects. Groups can contain collections of any type of object defined in a management pack and discovered and monitored by Ops Manager. Technically, every group is also a class (a special single-instance class called a singleton class). However, because groups can be used to create collections of subsets of the instances of a class for management pack tuning, they deserve special mention here.
- Health state: The current health of a monitored object (red, yellow, or green); sometimes simply called a state.
- Monitors: Monitors periodically assess the condition of specified objects, such as services and performance counters. As a result of the assessment, a monitor can change the health state of an object and can generate alerts. Only monitors understand health state (rules don't). As a result, monitors can be configured to automatically resolve an associated alert when the error condition improves and monitor state returns to healthy.

Ops Manager monitors include unit monitors, aggregate rollup monitors, and dependency rollup monitors.

- Unit monitors: These monitors are created to monitor specific aspects of applications, devices, and services, such as Windows events, services, and performance counters. They can also be used to monitor network devices through SNMP.
- Aggregate rollup monitors: You typically use an aggregate rollup monitor to group multiple like unit monitors into a single (aggregate) point, then use that monitor to set the health state (and generate an alert if desired). Aggregate monitors roll up the health of monitors beneath them according to a defined algorithm. These monitors use one of two algorithm elements: WorstOf or BestOf. With the WorstOf algorithm, if any unit monitor that's being aggregated is in a warning or error state, the aggregate rollup moves to a warning or error state. With the BestOf algorithm, if at least one of the unit monitors that's being aggregated remains in a healthy state, the aggregate monitor also remains in a healthy state.
- Dependency rollup monitors: These monitors roll up the health from an instance of another class linked by a hosting or containment relationship. For example, a monitor of this type rolls up the health of the IIS 7.0 Server Role class to the Windows Server 2008 instance hosting the role. Like aggregate rollup monitors, dependency rollup monitors support BestOf and WorstOf algorithms for health rollup. Dependency rollup monitors also support a third algorithm called WorstOfAPercentage. The monitor using this algorithm will transition to an unhealthy state when X percent of the monitors it aggregates are in an unhealthy state (where X is a percentage defined by the management pack author or later by Ops Manager administrators).

Additional management pack components include object discoveries, overrides, rules, tasks, views, and reports.

- Object discoveries: An object discovery is used to dynamically find objects and their properties on the network that need to be monitored. Object discoveries can use registry information, Windows Management Instrumentation (WMI) queries, SNMP, scripts (VBScript, JScript, or PowerShell), or managed code to identify applications running on a managed system. Sometimes simply called a discovery.
- Overrides: An override is an adjustment to the default settings of an object discovery, monitor, rule, or task. Overrides are how administrators tune management packs in Ops Manager.
- Rules: Rules collect data, such as performance and event information, generated by managed objects. Rules can be configured to generate alerts. However, rules don't affect the health state of the objects they monitor.
- Tasks: A task performs an administrative action on demand when an Ops Manager administrator selects the task. Some tasks run on the managed computer (console tasks), and some tasks run on the computer from which they were initiated (agent tasks).
- Views: Views display a particular aspect of monitoring settings. The displayed information in a view is the result of a query to the Ops Manager database. Ops Manager views include Alert, Event, Performance, State, Diagram, Dashboard, Task, and URL.
- Reports: Reports (delivered in many management packs) can display information about object availability, as well as event, alert, configuration, or performance data collected from the applications or devices monitored by the management pack.

Management Pack Design

A management pack's design determines how effective the management pack is in monitoring the application it targets. Although the experience level of the management pack author is important, it's equally important to involve a subject matter expert in the authoring process who truly understands the application to be monitored. A management pack created by someone who isn't familiar with the architecture and function of the target application won't result in an effective monitoring solution for that application.

How does Ops Manager dynamically discover and monitor applications? It all begins with the management pack design process, in which the authors define the application components of interest and how the components are related.

Service model. The combination of classes and their relationships within a management pack is referred to as the application's service model. Figure 1 shows the Windows Server 2008 Management Pack's service model.

The application components of interest for monitoring are defined as unique object types, or classes. An occurrence of a class discovered by Ops Manager is referred to as an instance of the class.

For example, the Windows Server 2008 Management Pack contains a class called Windows Server 2008 Computer, as Figure 1 shows. It also contains several classes representing the components of Windows Server 2008 Computer, such as Windows Server 2008 Operating System, Windows Server 2008 Logical Disk, Windows Server 2008 Physical Disk, and Windows Server 2008 Processor—object types of interest from a monitoring perspective.

However, it's not enough to describe and discover instances of these object classes. If every time we removed an Ops Manager agent from a Windows Server 2008 computer, its processor, disk, and other components remained in the Operations console, we'd have a serious administration headache. To dynamically discover, monitor, and (when necessary) remove instances, Ops Manager must also understand how the objects are related. To address this need, the management pack author creates relationships to describe the dependencies and interaction between object types. Objects can be connected through one of three types of relationships.

- Hosting relationship: For example, the Windows Server 2008 Management Pack contains a class called Windows Server 2008 Operating System. This class is said to be hosted by the Windows Computer class because the Windows OS can't exist without the computer that it runs on.
- Containment relationship: In this type of relationship, instances of the target class can exist even if they aren't part of a containment relationship. For example, a computer is contained within a computer group, but the computer can exist even if it isn't part of a computer group.
- Reference relationship: In this type of relationship, object types aren't dependent on each other. For example, a database could reference another database that it's replicating with. This is the least specific type of relationship (as well as the least often utilized by management pack authors).

After an application's object types (classes) and their relationships are defined (establishing the service model), the author(s) can add the logic to define how the health of one object type affects another related object type.

Health model. After the object types and their relationships are established, the author can define how the health of each object type affects the health of others in the management pack. An application's health model determines how the current state of an application or application component is represented in Ops Manager. This model shows how the health state of individual application components affects the overall health of the application as displayed in the Operations console.

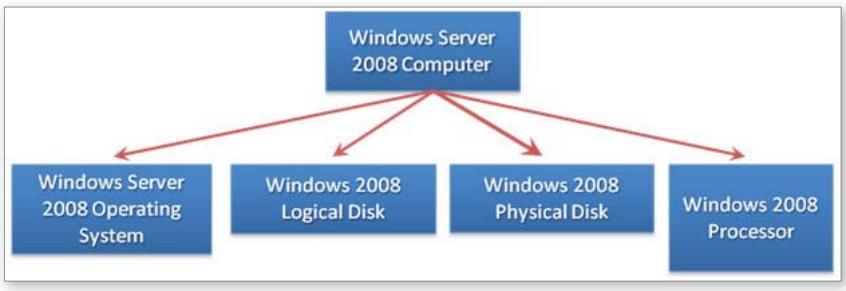

Figure 1: Windows Server 2008 Management Pack service model

The health of a hosted object class isn't automatically reflected (rolled up) in the health of the parent class—at least not in all cases. It's important to understand what's necessary to define the health rollup of the application components in a management pack.

Ops Manager health rollup requires two elements:

- Relationships: Health of the instances of one class can be rolled up to another class only if a hosting or containment relationship exists.
- Rollup monitors: A dependency rollup monitor must be created to roll up the health of one class to another.

You might wonder whether the health of a class *should* be rolled up to another class. This is a question the author must ask when designing a management pack's health model. The short answer is "it depends"—let's look at a couple of common examples to provide some perspective.

In the case of the Windows Server 2008 Management Pack, the health of the Windows Server 2008 Operating System class is rolled up to the Windows Server 2008 Computer class that hosts the OS. Because the primary function of the computer is to host the OS, it's logical to assume that if the OS is unhealthy, the computer won't be fully functional. The authors therefore included a dependency rollup monitor to transfer the health of the OS to the host computer.

In the case of the SQL Server 2008 Database Engine class (which represents a SQL Server instance), the SQL Server 2008 Database Engine class hosts the SQL Server 2008 Database class (SQL Server 2008 DB). Although these two classes share a hosting relationship (as in the previous example), the application health dynamics are unique. Even if one or more databases are unhealthy (e.g., offline, corrupt, suspect), this doesn't necessarily indicate a problem with the SQL Server database engine that hosts the database(s). In this case, the management pack authors chose not to roll up the health of the database class to its database engine host.

Discovering Applications and Components

Object discovery is the management pack element that's used to dynamically find objects and their properties that need to be monitored. Object discoveries can use registry information, WMI queries, SNMP, scripts, or managed code to identify applications running on a managed system. Application discovery is both a repetitive process and a progressive process in most cases. This is possible because every object discovery (and every rule, monitor, and task) has a target (class). The object discovery runs only on computers with one or more instances of the class targeted by the object discovery.

This concept is most easily explained by reviewing an example from an actual Microsoft management pack. Let's walk through the discovery process for the Windows Server 2008 IIS 7.0 management pack, from import of the management pack, transfer to the agent, loading of workflows, and the high-level discovery process. This example also illustrates the concept of progressive discovery, whereby layers of deeper discovery using scripts and other probes are performed only on systems on which the application is known to exist based on a lighter weight registry discovery.

1. First, the management pack is installed through the Import Management Packs wizard in the Operations console.
2. The management pack is delivered to all agents with at least one instance of the Windows Server class. (This action shows up as event 1201 in the Ops Manager event log on the monitored system.) In practical terms, every management pack is delivered to every properly functioning agent-managed system with at least one instance of the class targeted by a discovery, rule, monitor, etc., in the management pack. The seed discovery runs only on systems hosting an instance of the class targeted by the discovery.
3. Agents receive the management pack and load the workflows (discoveries, rules, monitors) into cache. (This action shows up as event 1210 in the Ops Manager event log on the monitored system.)
4. In the IIS 7.0 management pack, a discovery executes for all agents hosting an instance of the Windows Server 2008 class. This initial seed (root) discovery uses a simple registry probe to identify whether the IIS 7.0 role is installed on the system. Because discovery is a repetitive process, this discovery runs on a recurring basis every few hours.
5. An instance of the IIS 7.0 Server Role class is discovered on agents with IIS 7.0 installed. Likewise, if the IIS 7.0 role were uninstalled from the computer, the discovery would detect this and the IIS 7.0 Server Role instance, along with all the objects it hosts (e.g., websites, FTP sites, application pools), would be undiscovered (removed from) the management group automatically when the agent submitted an updated set of discovery data to the management group reporting this change.
6. Then, additional script-based discoveries targeting the IIS 7.0 Server Role class run only on those agents where an instance of the IIS 7.0 role was discovered, discovering specific IIS 7.0 components, such as the FTP server and SMTP server.
7. This process continues, with more discoveries targeting the Web, FTP, and SMTP Server classes to discover websites, FTP sites, and SMTP virtual servers hosted by the IIS 7.0 role on the server.

> Application discovery is both a repetitive process and a progressive process in most cases.

By using progressive layers of discovery in this manner, the management pack authors avoid running more detailed (and resource-intensive) scripts on systems where the target application or application components aren't installed, thus preventing unnecessary consumption of system resources.

On an ongoing basis, the agent runs all the applicable discoveries that it knows about when the health service starts up (e.g., when you restart the health service, or after a reboot). The agent sends this discovery data to the management server, then runs the discoveries again based on the interval specified on the object discovery—which might be as frequent as once per hour or as seldom as once per week. When the discovery runs, the agent inspects the discovery data it receives and compares it to the last discovery data it sent to the management server. If nothing changed since the last time discovery ran, the agent drops the discovery data and forwards nothing. If anything changed in the discovery data values, the agent resubmits the new data to the management server, which then submits the data to the operational database. The Root Management Server (RMS) detects the change and recalculates (and regenerates) the configuration. This action shows up on the RMS as event 21025 in the Ops Manager event log.

Event 21025 doesn't indicate a problem; it simply indicates that the System Center Management Configuration service (OMCFG, sometimes called the OpsMgr Config service) performed as expected. In this case, the service successfully regenerated the configuration file from the data in the operational database, writing it to the following file: \%ProgramFiles%\System Center Operations Manager 2007\Health Service State\Connector Configuration Cache*MGNAME*\OpsMgrConnector.Config.xml.

If changes in your management group are so frequent that the RMS is in a perpetual process of recalculating the configuration, you might encounter a condition known as config churn. For an in-depth explanation of config churn, including how to detect, troubleshoot, and prevent this undesirable condition, see the following blog posts: "Troubleshooting 21025 events—wrap up" by Daniele Grandini (nocentdocent.wordpress.com/2009/07/09/troubleshooting-21025-events-wrap-up) and "What is config churn?" by Kevin Holman (blogs.technet.com/b/kevinholman/archive/2009/10/05/what-is-config-churn.aspx).

Monitoring and Health State Calculation

After one or more instances of an application and its components are discovered (represented by class instances and

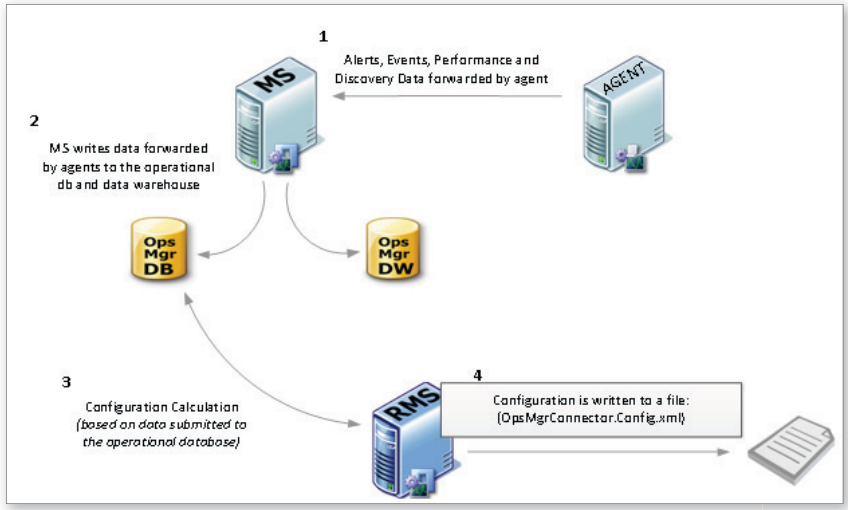

Figure 2: High-level data flow and process in configuration calculation

relationships), application monitoring begins. The rules and monitors contained within the management pack are loaded into cache by the agent as workflows. For example, in the management pack for IIS 7.0, unit monitors verify that Windows services related to SQL Server 2008 are running, confirm that databases are online and have adequate free space, and watch for error events in SQL Server and Windows event logs.

The agent runs these workflows and forwards this information to its designated management server. The management server writes this information to the operational database (Ops Manager), as well as to the data warehouse, if defined in the workflow. Data forwarded by the agent includes a number of different types, such as discovery, alert, event, and performance

The RMS then calculates the health state of monitored objects. After data is written to the Ops Manager database, the RMS watches and reads the instance space of the database and calculates the health state. As for configuration based on discovery data, calculation is controlled by the System Center Management Configuration service. Figure 2 illustrates the high-level flow of the configuration calculation process.

Tuning a Management Pack

Tuning a management pack involves using overrides to change the default settings. In sealed management packs, the management pack authors can control which discovery and monitoring parameters can be adjusted (tuned) by Ops Manager administrators by defining parameters as overridable parameters. Creating overrides is a relatively simple process, but knowing which overrides take precedence when competing or conflicting overrides exist can be confusing. However, you need to understand which overrides take precedence in order to avoid unnessesary problems.

The RMS determines which override takes precedence when multiple overrides apply to a workflow for a specific target, based on the following criteria:

- The first rule of overrides is that the most specific override takes precedence. For example, an override targeted to a group takes precedence over an override targeted to a class. Therefore, an override targeted to the Windows Server 2008 Computers group wins over an override targeted to the Windows Server 2008 class because the group, which contains a subset of the instances of the class, is the more specific target. As another example, an override targeted to an instance takes precedence over an override targeted to a group. Therefore, an override targeted to Server1 wins over an override targeted to the Windows Server 2008 Computers group because the instance is the more specific object.
- An enforced override takes precedence over all non-enforced overrides of the same type. For example, an enforced override targeted to a class wins over a non-enforced override targeted to a class.
- Overrides in unsealed management packs always take precedence over sealed overrides. This allows the administrator to better control settings in cases when an author has included overrides in a sealed management pack.
- Class overrides from hosted instances always take precedence over class overrides of the hosting instance.
- When all criteria for arbitration are exhausted, the RMS selects an override at random. It's important to never create so many like overrides that random selection is the only means by which Ops Manager can choose the effective override. One way to prevent this situation is to avoid creating group overrides for multiple groups that contain some of the same computers as members.

For a technical description of override precedence, see Jakub Olesky's "Overrides in SCOM 2007" at blogs.msdn.com/b/jakuboleksy/archive/2006/12/06/overrides-in-scom-2007.aspx. To learn more about Microsoft best practices for creating and storing overrides, see "Best practices to use when you configure overrides in System Center Operations Manager 2007" at support.microsoft.com/kb/943239.

Smooth Tuning

Tuning management packs can seem complicated. But knowing a management pack's design and function, understanding override precedence, and following Microsoft best practices for targeting and management pack tuning will go a long way in smoothing the tuning process.

InstantDoc ID 126053

Pete Zerger
(pzerger@akosts.com) is a consultant, blogger, author, speaker, and Microsoft MVP focusing on Microsoft System Center and virtualization technologies. He is a co-founder of systemcentercentral.com, a web community dedicated to System Center technologies.

Extending the Active Directory Schema

With proper planning, you'll have no fear

by Brian Desmond

Historically, both Active Directory (AD) administrators and IT managers have been fearful of extending the AD schema. Much of this fear stems from Microsoft documentation in the Windows 2000 era that made schema extensions appear to be dangerous and something best done with extreme caution. However, with a bit of planning and due diligence, extending your AD schema doesn't have to be something to fear.

The AD schema defines the structure of the data stored in the directory. Out of the box, AD supports many different types of objects (e.g., users) and attributes (e.g., first and last name). When the base schema that comes with AD doesn't lend itself well to data you need to store in the directory, you can extend the schema with custom objects and attributes.

Typically, the AD schema is extended for a number of reasons. For many organizations, the most common reason is the implementation of an application that requires a schema extension. Microsoft Exchange is a perfect example of this. Third-party software vendors also sometimes require schema extensions to support their application. Also quite common is extending the schema to support an internally developed application, or to provide a location to store proprietary data in AD.

Data Storage Options

The first things you should evaluate when considering a schema extension, particularly for an inhouse application, is whether or not the data is appropriate for storing in AD. Particularly well suited to AD is data that is relatively static (i.e., it doesn't change often), that is necessary across your organization (because it will replicate across domains), and that isn't highly sensitive (e.g., you shouldn't store birth dates, social security numbers, and so on in AD).

If you have data that doesn't match this criteria but still needs to be in an LDAP directory, a second option might be a good fit for you. AD Lightweight Directory Services (AD LDS, formerly ADAM) is a standalone version of AD that can run as a service on a member server (or domain controller—DC) and be queried via LDAP, just like AD. Rather than being constrained by the need to place AD DCs to support enterprise authentication and application requirements, you can tightly control who can read data and where that data is replicated by only placing AD LDS instances in appropriate locations.

Data Storage Primitives

To understand the AD schema, you need to know two key terms: class and attribute. Everything in AD, including the schema itself, is defined in terms of classes and attributes. Classes are types of data you want to store. For example, *user* is a class in AD, as is *computer*. Attributes are properties of classes. The user class has a first-name attribute (givenName) and a last-name (sn) attribute. The computer class has an OS attribute. The AD schema is defined in terms of two classes: classSchema for classes and attributeSchema for attributes.

If you're familiar with a typical database, another way to think of this is that classes are analogous to tables in the database, and attributes are analogous to columns within a table. Note, however, that you shouldn't get the impression that this is how the AD database (the directory information tree—DIT) is structured, because it's actually quite different.

When thinking about storing a new type of data in AD, you need to think about how the data maps to classes and attributes. For most common extensions, adding an attribute to an existing class (e.g., user or group) is sufficient. When you simply need to store a new piece of data about an existing type of object (such as users), you should first determine whether an existing attribute in AD is appropriate. The schema contains thousands of attributes and most of them are unused. So, for example, if you wanted to store a user's mail stop, you might consider the physicalDeliveryOfficeName attribute of users.

Repurposing an attribute for something other than its intended use is a bad idea. Consider the scenario in which you repurpose an attribute for something other than its intended use and then you buy an application that uses the attribute for its original purpose. You need to do double the work because you have to reconfigure the existing application using the attribute and then move the data. In general, it's always safer to add a custom attribute than to take this risk.

But sometimes you need to think in terms of classes; in a couple of scenarios, adding new classes to the schema makes more sense than using attributes. The first scenario is when you need to track a totally new type of data in the directory. If, for example, you wanted to keep track of the company cars in AD, it would probably make sense to define a new *car* class in the schema. Another scenario is when you need to do a one-to-many mapping.

Microsoft Exchange Server 2010 provides a perfect example of this. Each mobile device a user has synchronized to Exchange using ActiveSync is stored as an instance of a special object class (msExchActiveSyncDevice) in the directory. These mobile device objects are stored as child objects under the user who owns the device. This design permits the mapping of a significant number of attributes (for each device) to a single user.

Schema Extension Inputs

To create a custom schema extension, you need to gather a number of key inputs before you can implement your custom attribute or class in a development environment. Many of these inputs are required to be globally unique, so it's important to do the necessary prerequisite work before proceeding. Cutting corners on this prep work is how schema extensions become dangerous.

The first thing you need to decide on is the name of your class or attribute. The most important part of the name is the prefix. Because attributes and classes need to have unique names in your schema (and your customers' schemas, if you're selling an application), adding a prefix helps ensure that your ID attribute doesn't conflict with someone else's ID attribute. Typically, you'd use an abbreviated form of your company name for the prefix. For example, I use bdcLLC as the prefix for attributes that my company (Brian Desmond Consulting, LLC) creates. You might use abcCorp if you are ABC Corporation. Just think about the uniqueness of your prefix because no overall registry of prefixes exists. If you work for a company with a very common name, or an abbreviated name, think about how to make it unique.

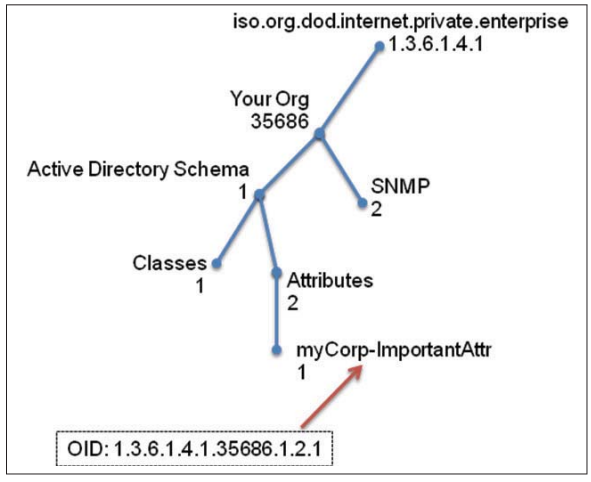

Figure 1: OID tree diagram

After you decide on a name, think about the Object Identifier (OID) you'll assign to your attribute or class. OIDs are an additional component that needs to be globally unique. AD (more generically, LDAP) isn't the only thing that uses OIDs for identifiers, so the Internet Assigned Numbers Authority (IANA) assigns unique OID trees to organizations upon request. Requesting a Private Enterprise Number, which is the portion of the OID tree unique to your organization, takes about 10 minutes and is free. You'll want to do this before you start creating custom schema extensions. You can request a Private Enterprise Number at www.iana.org/cgi-bin/assignments.pl.

After you have a Private Enterprise Number, you can create a seemingly infinite number of unique OIDs and organize them. Figure 1 shows a diagram of the OID tree for my company's Private Enterprise Number. Because you create OIDs by appending branches to the tree, many organizations first begin by creating an AD Schema branch (1.3.6.1.4.1.35686.1 in Figure 1), and then under that branch they create a class branch and an attribute branch. Under each of these branches OIDs will be allocated for each new attribute or class. In Figure 1, I allocated an OID (1.3.6.1.4.1.35686.1.2.1) for a custom attribute, myCorp-ImportantAttr. It's extremely important that you devise an internal tracking solution (such as an Excel spreadsheet or SharePoint list) to ensure that OIDs are always uniquely allocated.

Microsoft provides a script you can run that will generate a random OID, but you have no guarantee it will be unique. Best practice is to request a unique assigned branch from IANA and to use that for your schema extensions. Given how easy this process is, you'll never have reason to use Microsoft's OID generation script.

The remaining two inputs you need to decide on are specific to attributes and depend on the type of attribute that you want to implement. Linked attributes, which are extremely useful, are used for storing relationships between objects in AD. They are stored as pointers in the

AD database so that the relationships are always up-to-date in relation to the whereabouts of the object in the forest. Two common examples of linked attributes are group membership (member and memberOf) and manager/employee relationships (manager/directReports). Discussions of linked attributes often include the concepts of forward and backward links. The forward link is the editable portion of the linked attribute relationship. For example, with group membership, the member attribute on the group is the forward link; the memberOf attribute on the user is the backward link. When editing a group's membership, the modification must be made to the member attribute (the forward link) rather than the member object's memberOf attribute (the backward link).

To define linked attributes in AD, you need to define two attributes (the forward and backward links) and attach a link identifier (linkID) to each of these attributes. Link IDs need to be unique within the forest, and because other applications that might extend your schema need to use link IDs too, you'll want yours to be globally unique. Microsoft used to have a process for issuing link IDs to organizations, but in Windows Server 2003 the company replaced that process with a special indicator to AD that allows AD to generate unique link IDs when you extend the schema with a linked attribute pair.

AD expects link IDs to be sequential numbers. Specifically, AD expects that the forward link attribute is an even number, and the next sequential number is assigned to the backward link attribute. For example, with member and memberOf (group membership), the link ID for member is 4, and the link ID for memberOf is 5. If you need to support Windows 2000 forests with your schema extension, you need to continue defining static link IDs in the manner described here. Otherwise, you should use the auto link ID process introduced in Windows Server 2003. To use the auto link ID process, you need to follow the steps below when you define your schema extension. When you build your schema extension, as discussed later in this article, you'll need these steps to construct the linked attributes—if you're using linked attributes as part of your extension.

1. Create the forward link first, using a link ID of 1.2.840.113556.1.2.50. Note that although this link ID value is an OID, Microsoft simply reserved this OID value for the special purpose of creating an auto link ID.
2. Reload the schema cache.
3. Create the backward link attribute, using a link ID of the name of the forward link attribute.
4. Reload the schema cache.

The second item that's unique to attributes, and is also optional, is the MAPI ID. MAPI IDs are specific to Exchange Server. If you don't have Exchange or your attribute doesn't need to be surfaced in the Global Address List (GAL), you can skip this section. MAPI IDs are used to display attributes on one of the property pages in the address book, such as the user's general details template that Figure 2 shows. If, for example, you want to display employee classification (e.g., full-time employee or contractor) in the GAL, you need to assign your attribute for this as a MAPI ID. After you assign a MAPI ID to an attribute, you can use the Exchange Details Templates Editor to add that attribute's data to the view provided in the GAL inside Microsoft Office Outlook.

MAPI IDs must be unique, much like OIDs and link IDs. In the past you had no way to generate unique MAPI IDs, so these IDs were always a sticky point in the realm of schema extensions. Fortunately, Windows Server 2008 introduced a process to automatically generate unique MAPI IDs within the directory to reduce the risk of duplicate MAPI IDs. To use this functionality, assign a value of 1.2.840.113556.1.2.49 to the MAPI ID attribute when you create the attribute. AD will generate a unique MAPI ID for the attribute after the schema cache reloads. Note that although this value is an OID, it's reserved within AD for indicating automatic MAPI ID generation, much like automatic link ID generation discussed earlier.

To summarize, you must consider three crucial inputs when planning a schema extension. The first is the name of the class or attribute; the second is the unique prefix that you'll assign to all of your classes and attributes; the third is the OID. You need to request a unique branch of OIDs from IANA to generate your OID. If you're going to create a linked attribute pair, you'll need a unique pair of link IDs. And if you're going to surface your attribute in the Exchange GAL, you'll need to use a unique MAPI ID. In the case of both link IDs and MAPI IDs, using the automatic generation process inside AD is much better than using static values.

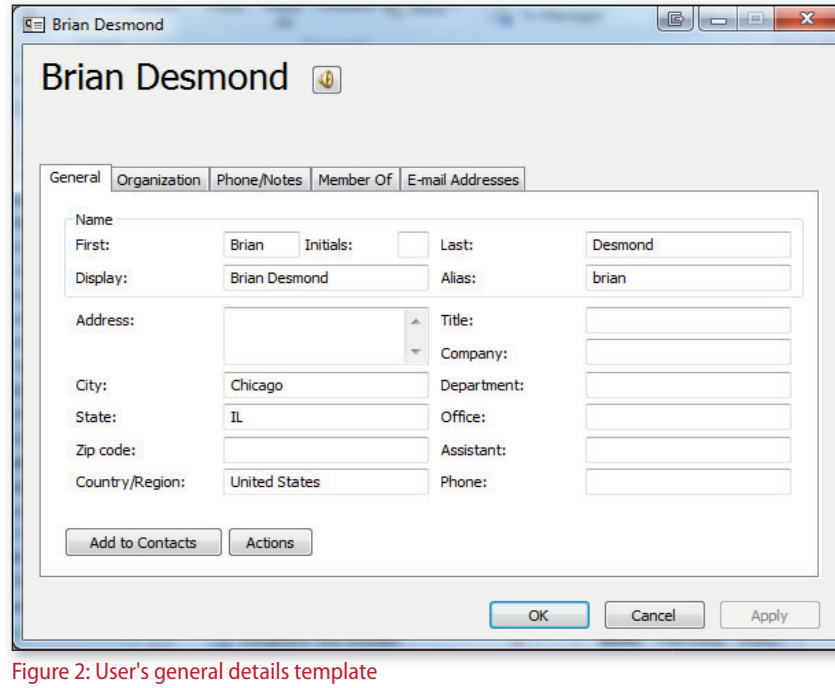

Figure 2: User's general details template

```
dn: CN=bdcllcShoeSize,CN=Schema,CN=Configuration,DC=X
changetype: add
objectClass: top
objectClass: attributeSchema
cn: sfsuLiveServiceEntitlements
attributeID: 1.3.6.1.4.1.35686.100.1.1.2
attributeSyntax: 2.5.5.12
isSingleValued: FALSE
showInAdvancedViewOnly: TRUE
adminDisplayName: bdcllcShoeSize
adminDescription:
 Stores a user's shoe size
oMSyntax: 64
searchFlags: 1
lDAPDisplayName: bdcllcShoeSize
name: bdcllcShoeSize
schemaIDGUID:: Js+e3rEsAUWMazlPm5hb6w==
isMemberOfPartialAttributeSet: TRUE
```

Figure 3: An example LDIF entry

Implementation Planning

When implementing a custom schema extension or extending your schema with a vendor's attributes and classes, you need to take some basic planning steps to protect the integrity of your AD forest. The first step is testing your schema extension.

If you're creating a custom schema extension, use a disposable development environment to create your schema extension. AD LDS is available as a free download for Windows XP and Windows 7 workstations. You can create an AD LDS instance on your workstation, design your schema extension in an isolated environment, and then export that extension for import in a test AD forest. AD LDS's schema is compatible with AD, so you can use LDIFDE for export. After you develop your schema extension, you can import it in your test AD forest and ensure the import is successful and that no key applications are affected. With regard to AD, you should plan to check for the import being successful and replication continuing to succeed within your test environment.

If you choose to test the schema extension with a test AD forest, it should have a schema that matches the production forest so that your testing is comprehensive. You can use the AD Schema Analyzer tool (included with AD LDS) to identify schema differences between two AD forests. The TechNet article "Export, Compare, and Synchronize Active Directory Schemas" (technet.microsoft.com/en-us/magazine/2009.04.schema.aspx) discusses how to import and export schema extensions, as well as how to use the AD Schema Analyzer tool. Note that there may be some differences depending on service packs and versions of Windows when you compare schemas, such as differences in attribute indexing and tombstone preservation.

In the case of schema extensions you didn't create (such as those bundled with a commercial application), you need to ensure that there's nothing suspect in the changes the vendor is planning to make. All the inputs we discussed in the previous section are critical to examine, as well as a few other things. The following list details the key variables you need to check:

- Delivered as an LDIF file (or series of LDIF files)
- Properly prefixed attributes
- Registered OIDs
- Registered/automatically generated link IDs
- Automatically generated MAPI IDs

LDIF files are an industry standard; all schema extensions you receive should be delivered in this format. It's permissible for applications to provide a custom import mechanism rather than requiring you to use LDIFDE to import the schema extension. But if the extension isn't delivered in this format, you should question the validity of it, as well as the practices of the vendor that created it. Figure 3 shows a sample LDIF entry that creates an attribute in the AD schema for storing a user's shoe size. You should note the following in this sample schema extension:

- The attribute is prefixed with the name of the vendor's company (Brian Desmond Consulting, LLC: bdcllc)
- The attribute has a unique OID issued under a Private Enterprise Number registered to the vendor
- The attribute is indexed (searchFlags: 1) and is available in the Global Catalog (isMemberOfPartialAttributeSet: TRUE)

You also need to ensure that an attribute's availability in the Global Catalog (Partial Attribute Set—PAS) is appropriate and that the indexes created on an attribute are appropriate if the attribute is going to be used in LDAP search filters. Also, it's wise to make sure that the data to be stored in the attribute is sensible for AD in the context of the restrictions and recommendations discussed earlier.

After you test your schema extension and you're ready to implement it in production, you'll want to plan for an appropriate time to do this. In general, it's perfectly feasible to make this change during business hours. You should expect some measurable increased CPU utilization on your schema master, as well as some negligible increased CPU utilization on DCs as they replicate the change. In large environments, you might also see transient suspensions of replication between DCs for four- to six-hour periods if you add attributes to the PAS. These suspensions will come with errors indicative of lingering object problems, but you can often ignore them and they will go away. If DCs remain quarantined from replication for an extended period of time, you should begin troubleshooting.

Following Through

Extending your AD schema isn't dangerous or something to be feared, if you take some basic steps. When planning new schema extensions, as well as evaluating custom attributes and classes from third-party vendors, look at the identifying information unique to each class or attribute and make sure that it's truly globally unique.

After you evaluate the integrity of the proposed extension, import it to a representative test environment and ensure that the test environment and key applications continue to function. Then you can import the schema extension into your production environment and begin taking advantage of it.

InstantDoc ID 126022

Brian Desmond
(brian@briandesmond.com) is a Directory Services MVP and senior consultant for Moran Technology Consulting in Chicago. Brian is author of *Active Directory*, 4th edition (O'Reilly), and blogs at www.briandesmond.com.

Exchange 2010 MRM

How to Modify and Reduce Help Desk Calls About Retention Policies

Mastering messaging records management

by Tony Redmond

Instead of using managed folders like its predecessor, Exchange Server 2010's messaging records management (MRM) system uses retention tags and policies to provide a way for users to control what's in their mailboxes and retain those messages and attachments that are required business records. In "Exchange 2010 MRM: Implementing New Retention Policies" (October 2010, InstantDoc ID 125359), I showed you how to design, create, and apply retention tags and policies using the Exchange Management Shell (EMS). Now I'll show you how to modify, remove, and customize retention policies. I'll also discuss how to help users understand them so you can reduce the number of Help desk calls.

Modifying Retention Policies

Policies often need to evolve over time, which means you'll need to add or remove tags. Exchange 2010 SP1 makes this task a lot easier because Microsoft has added a GUI for managing retention policies and tags to the Exchange Management Console (EMC). However, it's still good to know how to do the work in EMS, as this is the only way that you can work with retention policies and tags on an Exchange 2010 RTM server.

Adding tags. You can add new tags to an existing policy with the Set-RetentionPolicy cmdlet. To add a new tag, you must provide it along with the list of existing tags using the cmdlet's -RetentionPolicy TagLinks parameter. It's not sufficient to merely specify the new tag on its own, as this will update the policy to include only the new tag.

You can use two approaches to include a new tag in a retention policy. The first approach is to manually type in the complete list of tags. This approach works best for simple policies that include only a few tags, such as

```
Set-RetentionPolicy
  -Identity 'Audit Department'
  -RetentionPolicyTagLinks 'Audit-Inbox', 'Audit-SentItems'
Get-RetentionPolicy
  -Identity 'Audit Department'
```

(Although these commands wrap here, you'd enter each command on one line in EMS. The same holds true for the other commands that wrap.)

The second approach is to let Windows PowerShell retrieve the current list of tags, then add the new tag to that retrieved list. This approach works best when you have complex policies with many tags (e.g., six or more) and the potential exists that you might forget to input one of them.

Listing 1 demonstrates the second approach. This code first uses Get-RetentionPolicy to retrieve the existing tags, putting them in the $TagList variable. It also puts the new tag in the $NewTag variable. Next, the code uses the += operator to concatenate the new tag to the retrieved tags and assigns the

Listing 1: Code That Demonstrates the Second Approach to Adding a Tag to an Existing Policy

```
$TagList = (Get-RetentionPolicy `
  -Identity 'Management Retention Policy').RetentionPolicyTagLinks
$NewTag = Get-RetentionPolicyTag -Identity 'Manager-New-ArchivePolicy')
$TagList += $NewTag
Set-RetentionPolicy -Identity 'Management Retention Policy' `
  -RetentionPolicyTagLinks $TagList
Get-RetentionPolicy -Identity 'Management RetentionPolicy' |
  Select Name, RetentionPolicyTagLinks
```

resulting list back to the $TagList variable. Finally, it uses Set-RetentionPolicy to write the updated set of tags in $TagList to the policy. As you can see, the second approach requires a little more typing, but it has the advantage of absolutely guaranteeing that all the existing tags are preserved.

To show you how easy it is to do the same task with the new GUI in Exchange 2010 SP1's EMC, here are the steps you'd follow:

1. Go to the Mailbox node in Organization Configuration.
2. Select the Retention Policies tab.
3. Double-click the retention policy you want to modify.
4. Update the tags assigned to the policy as shown in Figure 1, then click OK.

Removing tags. To remove a tag from a policy, you have to write a replacement list into the policy using the first approach I described. Or, if you're using SP1, you just need to select the tag to remove, click the "X" icon, then click OK.

If you remove a tag from a policy, users covered by the policy can't apply the tag to any items in their mailbox. However, existing items that have been stamped with the tag continue in place and will be processed by the Managed Folder Assistant (MFA). This situation continues until the user explicitly assigns a replacement tag to an item or you remove the tag from Active Directory (AD) using the Remove-RetentionPolicyTag cmdlet. When you do the latter, the MFA will remove the deleted tag from items the next time it runs.

Removing Retention Policies

When you want to remove a retention policy from a mailbox, you simply set the policy to $Null. Here's an example of the command you'd use

```
Set-Mailbox
  -Identity 'JSmith'
  -RetentionPolicy $Null
```

When you want to remove the retention policy from the organization (i.e., AD), you use the Remove-RetentionPolicy cmdlet in a command such as

```
Remove-RetentionPolicy
  -Identity 'Retention
  Policy - PR Department'
```

Removing a retention policy has the effect of removing the policy from any mailboxes to which it's currently applied. If any mailboxes are associated with the policy, Exchange will prompt you to confirm its removal. If you proceed, Exchange removes the reference to the now-deleted policy from the mailboxes. Because Exchange won't know which retention policy should replace the one that was just removed, no policy is applied. Therefore, locating the mailboxes that a retention policy is applied to is a proactive step that you should take before you remove the policy. You can scan mailboxes to discover where a retention policy is applied with a command like this:

```
Get-Mailbox |
  Where {$_.RetentionPolicy
  -eq 'Retention Policy - PR
  Department'} |
  Select Name
```

You can easily adapt this command so that it applies a new policy to the mailboxes after it finds them. To do so, you use the Set-Mailbox cmdlet in a command such as

```
Get-Mailbox |
  Where {$_.RetentionPolicy
  -eq 'Retention Policy - PR
  Department'} |
  Set-Mailbox
  -RetentionPolicy
  "Public Relations Policy"
```

You'd run this command before running the command to remove the old retention policy.

Customizing Retention Policies for Specific Mailboxes

You can tailor the retention policy for a specific user by assigning personal tags on a per-mailbox basis. A personal tag is one that a user can assign to a folder or an individual item anywhere in his or her mailbox. This can only be done if a retention policy already applies to the user's mailbox. For example, suppose you want to assign a new personal tag that allows the user to mark an item to be moved into the archive after a year. This user's mailbox doesn't have any other personal tags, so you'd use the Set-RetentionPolicyTag cmdlet like this

```
Set-RetentionPolicyTag -Mailbox JSmith
  -OptionalInMailbox
  'Personal-Move-Archive'
```

Exchange adds the personal tag to the set of tags in the retention policy and makes the expanded set available the next time that the user connects. You can't assign a default policy tag (DPT) or a retention policy tag (RPT) to a mailbox in this manner. These tags can only be assigned to a mailbox through a retention policy. See "Exchange 2010 MRM: Implementing New Retention Policies" for detailed descriptions of DPTs, RPTs, and personal tags.

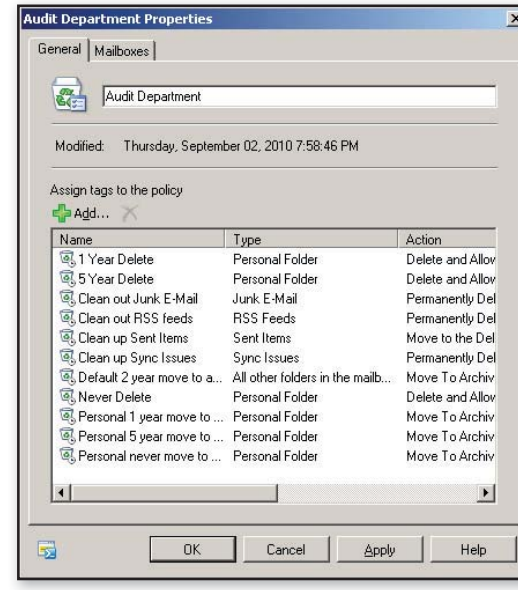

Figure 1: Modifying the tags assigned to a retention policy in Exchange 2010 SP1

November 2010

The Essential Guide to Microsoft Virtualization

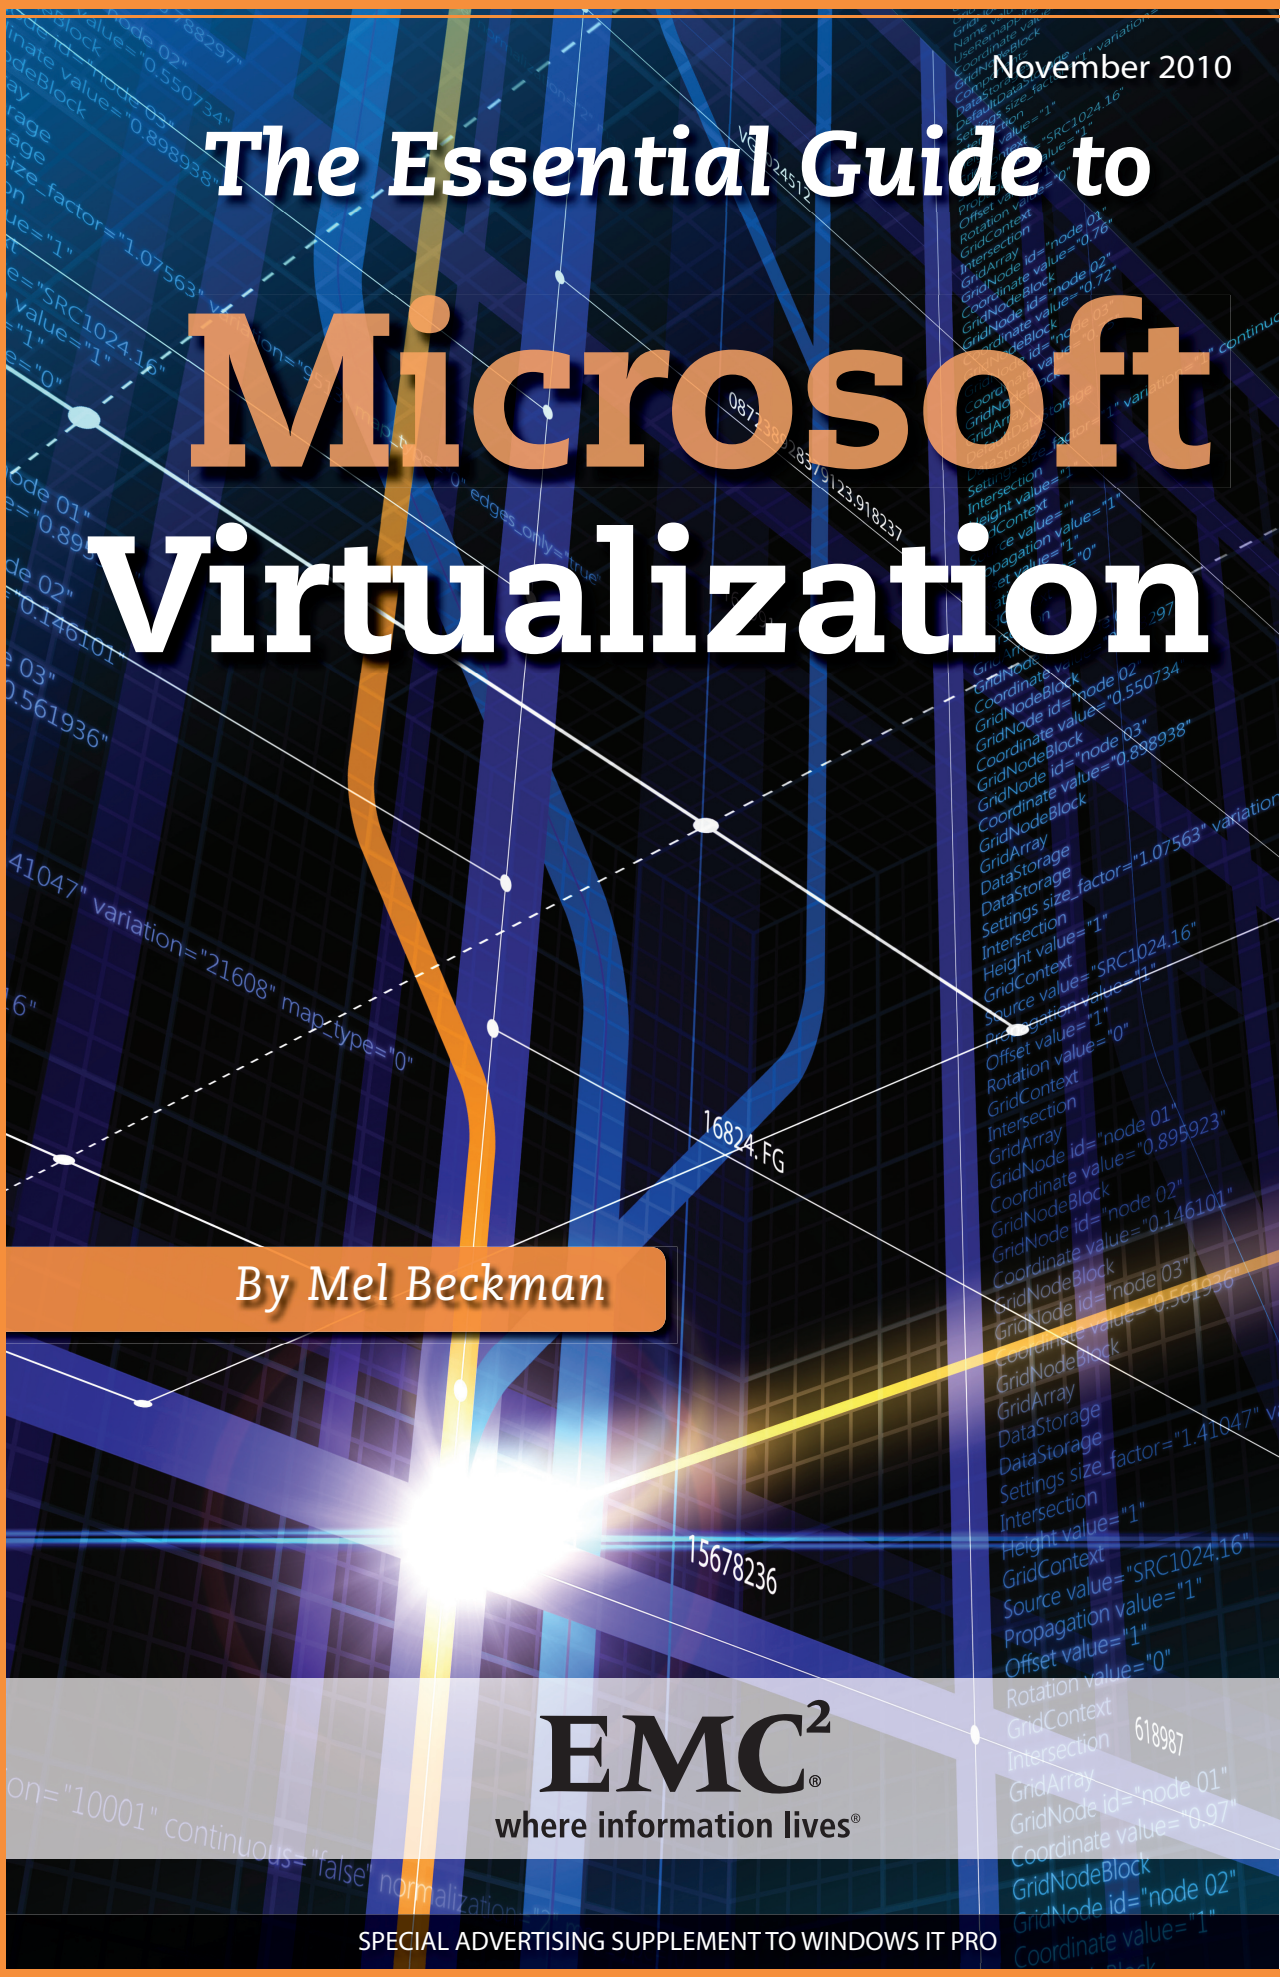

By Mel Beckman

EMC²
where information lives®

SPECIAL ADVERTISING SUPPLEMENT TO WINDOWS IT PRO

The modern enterprise has many missions: customer relationship management, enterprise resource planning, business intelligence, governance compliance, information security, to name just a few. All of these missions are critical. Not surprisingly, in the multi-mission enterprise, demand for IT services—particularly mission-specific application servers—is rapidly increasing. And the one key component underlying the entire virtual enterprise is centrally managed storage, which provides the substrate upon which virtual machines are built and operated.

The role of IT is changing in response to this new multi-mission orientation. Where once IT supplied both technical and application expertise to enterprise departments, today IT delivers services to mission managers that maintain application expertise within their organizations. IT, in essence, delivers a set of consistent services from which mission managers may draw in support of their applications. That set of services is very similar to the portfolio of modern public cloud infrastructures: server hosting, secure networking, application backup and restoral, business continuity, and disaster recovery. A private, rather than public, cloud.

Fortunately, server consolidation through virtualization gives enterprise IT powerful tools for meeting this demand, supporting more applications with fewer resources. Virtualization brings economies of scale through better CPU utilization, and by sharing storage across cost-effective arrays interconnected via a Storage Area Network (SAN), effectively virtualization storage. Storage virtualization, in turn, improves ROI on one of the largest cost components in any data center. Ultimately the combination of virtualized servers and storage enable the ability to deliver services through a private enterprise cloud diversely located in widespread geographic regions across thousands of miles.

Microsoft's Hyper-V virtualization can dramatically improve IT cost containment while enabling new capabilities and improving agility as the enterprise proceeds on its journey to the private cloud. Microsoft's cohesive computing infrastructure lets IT professionals smoothly scale up its infrastructure to support more, and larger, enterprise missions. Hyper-V virtualization lowers costs by letting IT exploit faster processors, networks, and storage, while simultaneously lowering labor expenses.

Implemented correctly, virtualization simplifies IT service delivery. Implemented incorrectly, virtualization leads to reduced reliability and a higher probability of catastrophic failure, ultimately neutralizing virtualization's cost advantages. To successfully deploy a new Hyper-V-centric infrastructure, you need to know about the risks any form of virtualization can bring: reduced reliability, runaway server growth, and slower application performance. Fortunately, EMC can mitigate these risks and make your virtualized infrastructure more agile by giving you the tools and expertise you need to roll out a reliable and secure Hyper-V infrastructure from the start.

Risks and Rewards

Consolidating multiple physical servers onto a single Hyper-V host immediately introduces challenges for the organization, such as mitigation of potentially reduced storage level performance and higher impact of catastrophic physical server failure: an unavoidable, and subtle, concentration of risk. These risks can be mitigated—by using High Availability (HA) techniques to address single server failure, and by delivering robust storage solutions in the case of addressing performance concerns. Hyper-V has built-in features, such as virtual machine snapshots, Cluster Shared Volumes, and Data Protection Manager that directly addresses virtualization risks.

Managing risk should be high on your IT infrastructure requirements planning checklist. New governance rules add items to that list, by requiring explicit data encryption in more places than ever, with stricter data retention rules and new data discovery requirements to achieve governance compliance. All of these increase the need for backup data tracking and auditing.

Risk management thus becomes one of the services IT offers, along with a reliable computing infrastructure. By supplying backup, disaster recovery, and management processes to mission managers within the enterprise, IT truly becomes a private cloud provider.

The key to managing risks while delivering effective private cloud services is tight control over resources such as CPU, storage, and networking. EMC products such as PowerPath Virtual Edition let you deploy encryption at critical security checkpoints to meet governance objectives, and enhance Hyper-V HA through improved storage management to provide business continuity. EMC NetWorker and Replication Manager work with Hyper-V's Volume Shadow Copy Service (VSS) as a requestor service within VMs to streamline backups and application restoral.

Identifying the Mission-Critical Mission

In order to start consolidating servers with Hyper-V, you must assess the applications and

Top 10 Reasons to Use EMC for Microsoft Hyper-V Virtualization

1 You want to accelerate the benefits and TCO of virtualization with Microsoft Hyper-V and you want to deliver responsive—yet predictable—results to your users.
You can leverage Microsoft virtualization technologies with the latest EMC storage technologies for high performance, automated failover and failback, simplified provisioning, and virtual storage provisioning.

2 You need the right information infrastructure for virtualized environments and Private Cloud Computing.
As more companies move to hyper-consolidated virtualized data centers and the private cloud, performance and simplified management are increasingly important. EMC Unified Storage systems are so simple and efficient that EMC will guarantee our storage is 20 percent more efficient than competitive platforms. Additionally, EMC FAST and FAST Cache can help you execute a cost effective tiering strategy without manual intervention.

3 You want to lower costs and enhance efficiency of your business applications.
With virtualization from EMC and Microsoft you can lower TCO and improve agility and flexibility for next-generation deployments of Microsoft Exchange Server, Microsoft SQL Server, Microsoft SharePoint Server, and other key business applications. EMC's innovative storage technologies can reduce TCO by intelligently utilizing Flash and SATA drives.

4 You need network storage platforms that fully support Microsoft Windows Server 2008 R2 Hyper-V.
EMC's industry-leading network storage platforms fully support Microsoft Windows Server 2008 R2 Hyper-V and help provide a scalable, highly available information infrastructure for virtualized environments.

5 You are concerned about data protection and rapid recovery in virtualized environments.
EMC offers a range of data protection and rapid recovery options for virtualized Microsoft environments. EMC's expertise in aligning technology options to business requirements helps ensure data protection for address-consolidated virtualized data centers. EMC NetWorker supports a wide range of backup and recovery, including backup to disk, data replication, continuous data protection, and deduplication.

6 You want to work with an industry leader in Microsoft virtualization.
EMC Consulting's deep knowledge provides expertise in assessing, planning, and implementing Microsoft virtualization technologies. In fact, EMC won the 2008 Microsoft Partner of the Year for Business Process and Integration Solutions with a solution using Microsoft virtualization technologies.

7 You need assurance that your virtualized infrastructure can be supported.
Deep integration testing through the EMC E-Lab and a long-standing engineering relationship with Microsoft enables EMC to provide ongoing product support for Microsoft Hyper-V.

8 You want to deploy virtualization quickly and get it right the first time.
EMC is integrating Microsoft Windows Server 2008 R2 Hyper-V in its EMC Proven Solution development process. This integration results in solutions that combine EMC and Microsoft technology and that have been tested and documented for a broad set of common workloads and use cases. As a result, joint EMC and Microsoft customers have the opportunity to more rapidly achieve the cost reduction, efficiency, agility, and flexibility of virtualized environments.

9 You need virtualization-specific services.
To help you optimize your infrastructure and leverage virtualization, EMC offers Microsoft Hyper-V and Microsoft System Center assessment, planning, and implementation services; design services for application virtualization with Microsoft Hyper-V, and Virtual Desktop Implementations (VDI).

10 You want to work with a vendor with the scale to handle the entire enterprise.
EMC Global Services has thousands of consultants and offers a broad portfolio of strategic consultation, planning, delivery, and support across the entire IT lifecycle—from envisioning through day-to-day operations. Additionally, EMC is a Gold Certified partner with 15 certifications—including virtualization—and is a 20-time Partner of the Year Winner.

Take the next step
To learn more about EMC, contact your local EMC sales representative or authorized value-added reseller, contact us at Microsoft@emc.com, or visit our website at www.emc.com/microsoftvirtualization.

servers to be consolidated to choose the best consolidation candidates. That assessment should include measuring current CPU and storage utilization, and estimating future resource requirements.

At this point you should also define disaster recovery and business continuity requirements: which applications require HA, and which can sustain an outage within specific recovery time and recovery point objectives. These requirements will influence both the capacity and performance requirements of your Hyper-V infrastructure.

Armed with that information—which you can collect using Microsoft's Assessment and Planning (MAP) toolkit for Hyper-V—you're ready to start shopping for host, network, and storage platforms that meet these initial requirements, as well as accommodate predicted growth. EMC has expertise and products that can help you achieve project success from the get-go. EMC consulting services delivers focused engineering talent, with experts specializing in every EMC virtual storage product line. For example, EMC specialists can help you deploy VPLEX in an enhanced Hyper-V Cluster Storage Volume (CSV) implementation that eliminates physical barriers within, across, and between data centers.

Management Made Effective

Most Microsoft management products, such as Hyper-V Manager and System Center Virtual Machine Manager, provide support for Microsoft Management Console (MMC) interfaces, providing a useful foundation for some basic Hyper-V management capabilities. But highly focused management tools are often better suited to utilize all the features of a particular hardware or software component, accessing features and controls not accessible via generic MMC interfaces. Focused management gives you tighter control over mission-critical storage functions that span applications and Hyper-V hosts.

The best-focused management tool sets deliver a cohesive management console that displays current performance information alongside established metrics and service level agreement (SLA) requirements. They also incorporate comprehensive logging and alert mechanisms to escalate problem situations smoothly, providing the performance history and trend data necessary to resolve performance issues before they become critical.

Two common infrastructure-wide situations that a good management tool addresses are storage I/O contention and storage growth planning. The former is a tactical situation, requiring fast action to mitigate I/O contention before SLA performance and availability thresholds are breached. The latter is strategic in nature, but essential to ensuring the private cloud scales without impacting future mission requirements.

Another common management process that often includes tasks falling outside the scope of a generic MMC console is provisioning virtual machines and virtual storage. You can often automate some of these out-of-scope processes using Windows PowerShell cmdlets, but much routine provisioning can be performed by task-oriented management tools that have visibility across both virtual host and virtual storage infrastructures.

Storage Planning, Provisioning, and Management

Server and storage virtualization go hand-in-hand, but the planning, provisioning, and management processes are quite different. For example, expanding storage capacity often requires more advance notice than expanding compute capacity, as storage touches multiple aspects of the data center: networking, existing workloads, backup and recovery. Planning processes must be strategic, taking into account future storage requirements based on measured growth trends. Provisioning and management are tactical processes, based on immediate workload requirements and service requests from users.

EMC also supplies powerful MMC-compatible management components, such as Virtual Storage Integrator (VSI) for Hyper-V Manager, a graphical interface with PowerShell integration that provides a mapping between Hyper-V virtual machine objects found in System Center Virtual Machine Manager (SCVMM) and EMC storage. In addition, EMC storage arrays support all forms of storage connectivity employed by Hyper-V, including iSCSI and pass-thru disks, which you can manage within SCVMM System Center Operations Manager (SCOM).

Resilience Models

Achieving business continuity requires selecting tradeoffs between application resilience and the cost of mitigating various kinds of risks. Risks fall into three geographic categories: local, metropolitan, and regional. You can mitigate local risks using traditional disaster recovery tools, such as tape and disk-to-disk backup, or through enhanced backup/recovery models. For example, Hyper-V VM snapshots provide the ability to capture an application's runtime state, then later restart the application at the snapshot's point in time.

But backup/restore techniques impose recovery delays that are not acceptable for many mission-critical applications. For these you need fast fail-over technologies, such as Hyper-V failover clusters and SAN mirroring. Hyper-V Cluster Shared Volumes are one way to deliver fast fail-over, but tend to have application dependencies that can become cumbersome to manage. SAN mirroring can reduce this management burden, by transparently replicating

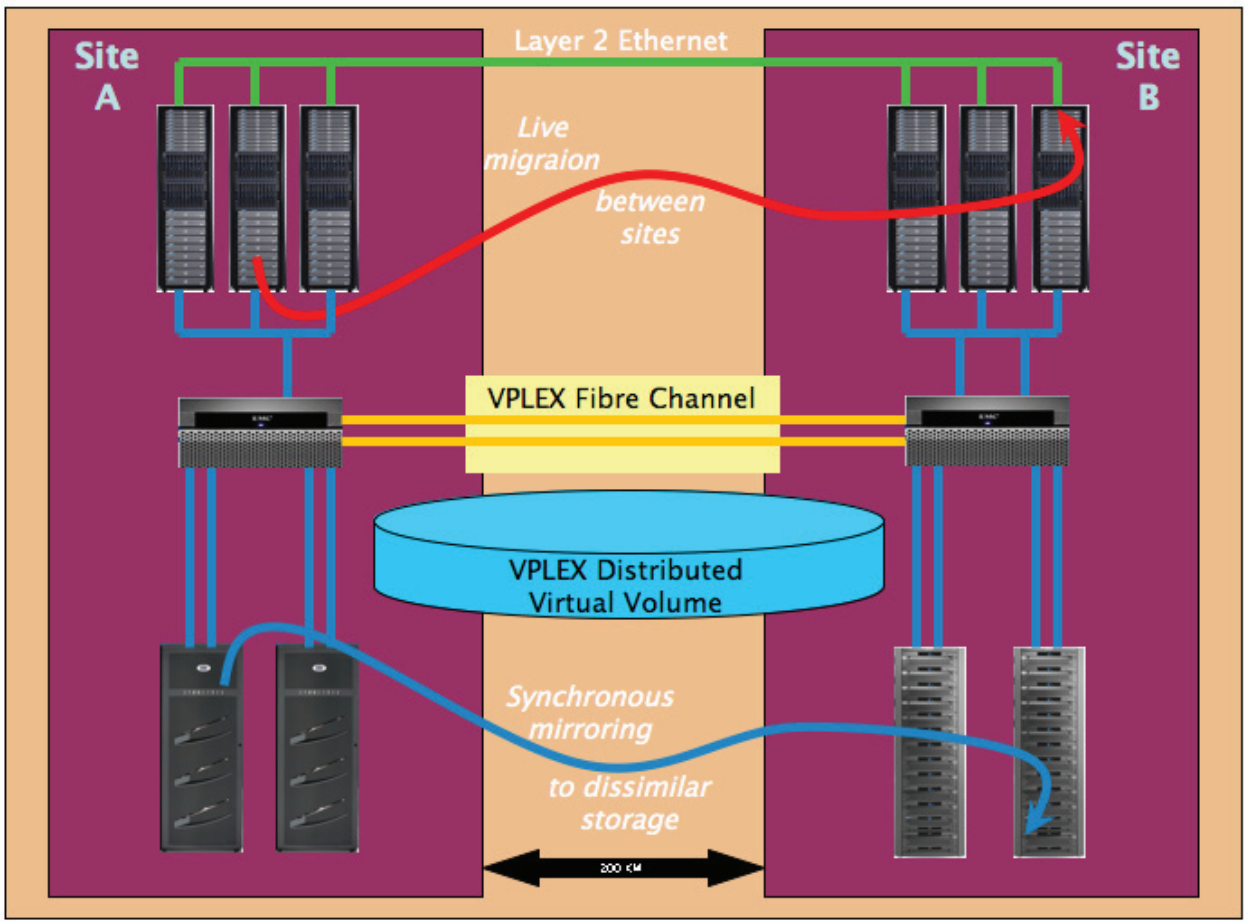

Figure 1: EMC VPLEX enables live migration to remote sites as well as synchronous replication for simplified High Availability

data between SAN instances deployed across a building, campus, or metropolitan area. As long as geographically dispersed SAN sites have low-latency interconnect, virtual workloads at one site can transparently migrate to a SAN instance at another site, protecting those workloads against the highly concentrated risk of catastrophic SAN failure.

EMC's VPLEX is an enterprise SAN federation platform that goes beyond typical mirroring solutions, with features such as advanced data caching to minimize latency and distributed cache coherence to extend synchronization distances. VPLEX can deliver virtual storage across networks stretching hundreds of kilometers, enabling extremely resilient application architectures and making the promise of private enterprise service clouds feasible today (Figure 1).

For SMBs, EMC CLARiiON MirrorView is an example of a highly available storage mirroring solution that replicates SANs across a datacenter or campus environment, maintaining byte-for-byte equivalence, either uni- or bi-directionally, with EMC CLARiiON SANs. EMC Symmetrix VMAX storage arrays provide a broader selection of storage replication capabilities at the high end, via Symmetrix Remote Data Facility (SRDF), which can mirror data across campus, metropolitan or even continental geographic separations. Tight integration for both CLARiiON and Symmetrix geographically dispersed solutions and Windows Failover clustering is provided by EMC Cluster Enabler.

For heterogenous storage environments, where different arrays are utilized in multiple locations, integration with Microsoft Failover Clusters is provided through EMC Cluster Enabler, combined with EMC RecoverPoint or RecoverPoint/SE. These solutions can also provide support for continuous remote replication between cluster notes, which can be geographically dispersed for seamless application recovery in Hyper-V environments.

EMC also supports the traditional backup mission with its NetWorker and Replication Manager families of centralized backup and data protection solutions. NetWorker automates backups and streamlines recoveries through a common management platform that supports disk-to-disk backup, deduplication, continuous data protection, and replication – across both physical and virtual machines.

Getting the Right Help

Storage system design and deployment requires more advanced skills than managing ongoing

operations, in the same way that engineering an aircraft requires more advanced skills than piloting one. Consultants and engineering services organizations can bring in the expertise you need for successful planning and deployment, and train your staff for efficient operations. You should evaluate potential services providers based on their experience in all sizes of projects, their competence with Hyper-V and other Microsoft products, and their library of best practices

EMC offers both design and implementation services, as well as pre-packaged solutions that leverage EMC's extensive experience with Microsoft Hyper-V. EMC's Information Infrastructure Solutions for Microsoft Virtualization offers a broad spectrum of hardware, software, and services support for Hyper-V and SCVMM, ensuring you obtain the agility, performance, and cost advantages virtualization can deliver. EMC specializes in wide-area, cross-site business continuity and disaster recovery, and is a Microsoft Geographically Dispersed Hyper-V Solutions partner.

With the viability of your entire virtualization deployment hinging on correct engineering, training, and ongoing operational management, EMC professional services are one of the most cost-effective means of ensuring a successful project.

Mel Beckman is a senior technical editor for Penton Media. He has built two regional Internet service providers and is currently president of Beckman Software Engineering, a technical consultancy specializing in large-scale, high-bandwidth networks. His past clients include Apple Computer, the City and County of Santa Barbara, DuPont Displays, IBM, Loral Federal Systems, United Airlines, the U.S. Department of Agriculture, and the U.S. Department of Energy. Mel has presented seminars on computer programming and network technology throughout the United States, Europe, and Asia.

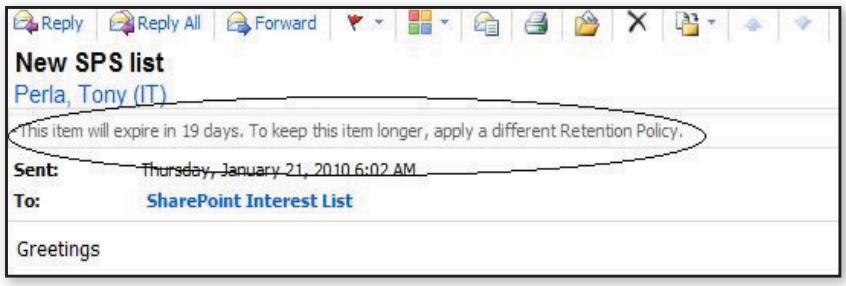

Figure 2: A warning about an approaching expiration date in OWA

Unfortunately, no cmdlet is available to retrieve a mailbox's assigned personal tags. (If you examine a mailbox with Get-Mailbox, it tells you if a retention policy is assigned, but nothing else.) Therefore, if you want to change the personal tags assigned to a mailbox, you have to write the complete list with Set-RetentionPolicyTag. For example, to add a personal tag to a policy that already has one assigned, you'd use a command such as

```
Set-RetentionPolicyTag
  -Mailbox JSmith
  -OptionalInMailbox
  'Personal-Move-Archive',
  'Personal-Keep-LongTime'
```

To remove all the personal tags from a mailbox, you set the list to $Null as follows

```
Set-RetentionPolicyTag
  -Mailbox JSmith
  -OptionalInMailbox $Null
```

Informing Users About Retention Policies

The first evidence users see that their mailbox has been assigned a retention policy is when they see warnings in message headers, which start appearing 30 days before an item expires. These warnings are visible when a message is opened or viewed in the preview pane. For example, Figure 2 shows a warning in Outlook Web App (OWA). It advises that a message has 19 days to go before it expires as the result of a retention policy tag placed on the Inbox. These warning messages can prompt a lot of Help desk calls if users aren't informed about the retention policies and the options they have when they get a warning message.

When users receive a warning message, they have two options: Let the item expire or apply a different tag to the item. If users decide to let an item expire, Exchange will perform whatever action is defined in the tag (e.g., move to Deleted Items, permanently delete) after the expiration date is reached.

Although the tags indicate the retention period (e.g., 6 months), they don't tell users what action will occur when the retention period expires. Tags have a variety of associated actions from permanently delete to merely warning that the retention period has expired. (An exception is an archive policy, as these tags always mean that an item will be moved to a personal archive when the retention period expires.) Outlook 2010 users can view the actions for the default tag on a folder by viewing folder properties, but this information isn't available in OWA. (Outlook 2010 and OWA are the only two clients that currently expose retention tags.) So, when you create tags, you might want to incorporate the action in its name. For example, a tag named "One year delete" gives a pretty good hint to users as to what might happen, whereas one named "RPT-MGR-Retain" does not.

If users choose to apply a different tag to an item in their Inbox, they can do so one of two ways. They can move it or assign a personal tag to it:

Move the item. By moving the item to a different folder, users are removing it from the influence of the RPT that applies to Inbox items. After it's moved, the item is governed by the DPT defined in the retention policy that applies to the mailbox (if it exists) or the policy applied to the folder and inherited by all items that are added to the folder. If neither condition exists, the item is left untagged and therefore won't be subject to processing by Exchange.

Assign a personal tag. Users can explicitly assign a personal tag to the item. They can choose from any of the personal tags defined in the retention policy that's applied to their mailbox. To do so, they just need to right-click an item and select the personal tag.

As you can see in Figure 3, the Retention Policy UI differs between Outlook (left) and OWA (right). Outlook's UI provides a richer set of options, but OWA's UI is less confusing for the novice user. Users will see the Retention Policy UI only when a policy is applied to their mailbox.

After a personal tag has been applied to an item, the item is no longer subject to the provisions of either the folder policy

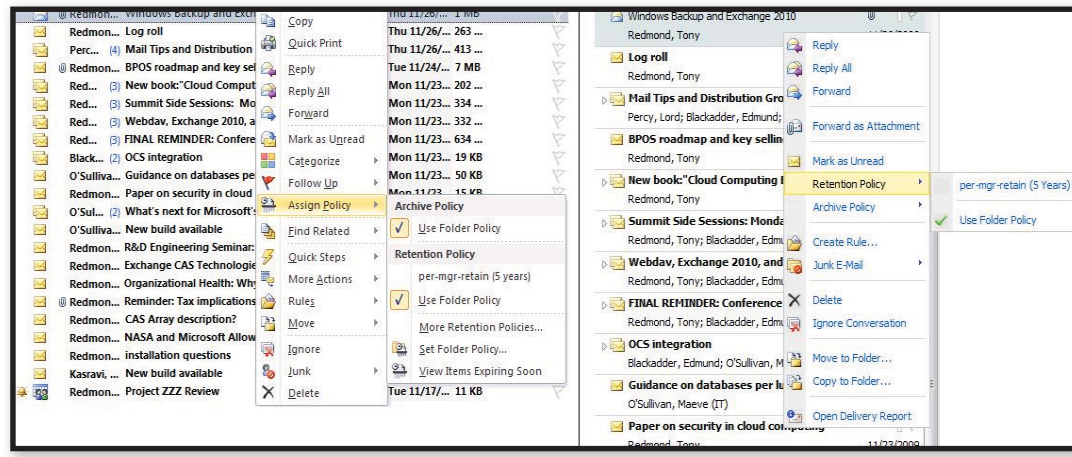

Figure 3: The Retention Policy UI in Outlook (left) and OWA (right)

Figure 4: The Archive Policy UI in OWA

or the default policy because an explicitly assigned personal tag always takes precedence over a tag placed on a folder. The personal tag remains with the item if the users move the item to another folder. If users want to impose a different retention tag on the item, they have to replace the existing personal tag with a different tag selected from the set available in the policy.

If a tag specifies "MoveToArchive" as its action, Outlook and OWA list it under the Archive Policy UI rather than the Retention Policy UI, as Figure 4 shows. These tags are under the Archive Policy UI because they can only be used with mailboxes that have personal archives.

Applying Advanced Settings to Facilitate Better Communication

Anything that facilitates better communication with users is likely to reduce Help desk calls, so you might want to have Outlook or OWA display additional information by setting some mailbox properties. For example, you can provide a message box like that in Figure 5 by setting two mailbox properties:

- RetentionComment, which you use to tell readers the retention policy that's applied to the mailbox
- RetentionUrl, which you use to provide a hotlink to information about the policy

You can use the Set-Mailbox cmdlet to set these properties. To create the message box in Figure 5, the command would look like

```
Set-Mailbox -Identity 'JSmith'
  -RetentionComment 'The management
  retention policy applies to this
  mailbox'
  -RetentionUrl
  '<a href="http://intranet.xyz.com/
  retentionpolicies.html">
  Retention Policy Information</a>
```

To see whether a mailbox has the Retention Comment and RetentionUrl properties set, you can run a command such as

```
Get-Mailbox -Identity 'JSmith' |
  Select Retent*
```

If your users speak different languages, you can provide localized versions of the retention tag's name and comment, which the Outlook or OWA client will display based on its language setting. An example of a command to provide localized text for a retention tag is

```
Set-RetentionPolicyTag
  -Identity 'Per-Mgr-Retain'
  -LocalizedRetentionPolicyTagName
  'en-IE: Irish version of tag name'
  'en-US: US version of tag name',
  -LocalizedComment
  'en-IE: Irish text comment'
  'en-US: US text comment',
```

Notice that the -LocalizedRetentionPolicyTagName and -LocalizedComment parameters have two values: one for clients whose language is set to en-IE (Ireland variant of English) and one for clients whose language is set to en-US (United States variant of English). When you set these properties, you need to specify the language identifier (e.g., en-IE, en-US) and the text (i.e., name and comment) in the appropriate language, using the format shown in the sample command. Be sure that you get accurate translations that clearly convey the meaning of the tag. Don't be tempted to cut corners and use school-quality translations or those that you might be able to procure free of charge from the Internet.

To see whether a retention tag has localized text, you can run a command such as

```
Get-RetentionPolicyTag
  -Identity 'Per-MGR-Retain' |
  Select Local*
```

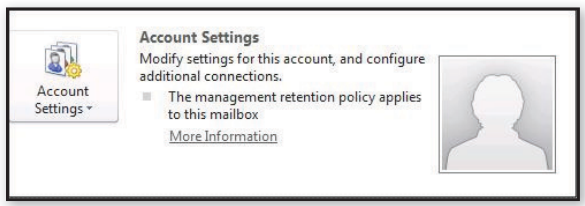

Figure 5: Message box telling a user which policy is applied to the mailbox

MRM: An Important Technology

Exchange 2010's implementation of retention policies establishes a comprehensive, automated, and user-friendly platform for administrators. Retention policies and tags will become an important part of many Exchange deployments, especially for companies that have to keep good records to meet legislative or regulatory requirements. In Exchange 2010, administrators need to use EMS to manage retention policies and tags, whereas in Exchange 2010 SP1, administrators can use EMS or the new GUI in EMC. I anticipate that the new GUI will make retention policies and tags more approachable for many administrators.

Retention policies and tags will become an important part of many Exchange deployments, especially for companies that have to keep good records to meet legislative or regulatory requirements.

InstantDoc ID 125919

Tony Redmond
(12knocksinna@gmail.com) is a contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 SP1 Inside Out* (Microsoft Press). His blog is at thoughtsofanidlemind.wordpress.com.

Scripting Shortcuts that Contain Unicode

How to work around WshShortcut's Unicode illiteracy

by Alex K. Angelopoulos

If your organization uses Unicode filenames, scripting shortcuts with Windows Script Host's (WSH's) WshShortcut object can be an exercise in frustration. What's particularly irritating is that this object allows you to write Unicode text to any of its properties without any fuss, but when you attempt to save the shortcut, WshShortcut either throws an error (if the shortcut's filename or path contains Unicode) or silently mangles the Unicode into garbage (if any other property contains Unicode).

Adding one final insult, WshShortcut returns the same generic error message—*unable to save file <filepath>*—for any refusal sent back by the Windows file system. The Unicode characters in the file path will be replaced with a question mark (?), making the error message even more confusing.

Fortunately, there are ways to work around these problems, which I'll explain. First, though, I'll shed some light on why WshShortcut exhibits these annoying behaviors.

Why WshShortcut Is Unicode Illiterate

The reason why WshShortcut accepts Unicode is easy to explain: It has to. Quoting from Microsoft's own COM interface design rules: "All string parameters in interface methods must be Unicode" [msdn.microsoft.com/en-us/library/ms692709(VS.85).aspx]. In WSH, all text written to a file (or another type of object) is represented internally as Unicode. (WSH and its scripting languages always use Unicode internally.) This pushes string conversion problems down to the level of the components that produce the output in a particular representation.

So why does WshShortcut throw errors and mangle content even though you're allowed to use Unicode in shortcuts? There's no official explanation for this, but it makes sense if you understand the historical context of WSH.

The WSH helper objects, including WshShell and its member objects (e.g., WshShortcut), were designed when Windows 98 and Windows 95 (both singularly lacking in Unicode file-system support) dominated the desktop. At that time, three simplifications were made:

- The designers wanted the WshShortcut object to function on all supported platforms. The easiest way for the same codebase to work identically on all systems is to use non-Unicode APIs, so that's what WshShortcut uses.
- Instead of inspecting the text it's given, WshShortcut simply assumes that it's ANSI text represented in Unicode. If you supply ANSI text to the object, there's no problem. When ANSI text is represented in Unicode, the first byte in each pair of bytes is empty, and WshShortcut's technique of chopping the first byte of each byte pair is no problem. However, when you supply Unicode characters, you now only have the representation of the last half of each character. This is essentially like chopping off the first half of the numbers in a street address. The address *8922 North Main* becomes *22 North Main*, which is not only wrong but also might not even correspond to a real address. Thus, when text that needs Unicode representation gets treated in this manner, you get garbage.

Listing 1: UnicodeShortcut1.vbs

```
UnicodeShortcut1.vbs *
1  Set WshShell = CreateObject("WScript.Shell")
2  Set ShApp = CreateObject("Shell.Application")
3  DesktopPath = ShApp.Namespace(0).Self.Path
4  bpmf = "ㄅㄆㄇㄈ"
5  ' Unicode paths like next line will fail.
6  ' Failure happens when saving the link.
7  'bpmfName = "ㄅㄆㄇㄈ.lnk"
8  bpmfName = "bpmf.lnk"
9  shortcutPath = DesktopPath & "\" & bpmfName
10
11 Set lnk = WshShell.CreateShortcut(shortcutPath)
12 lnk.TargetPath = "C:\Windows\System32\Notepad.exe"
13
14 ' Unicode internally just writes garbage
15 'lnk.Description = bpmf
16
17 lnk.Save()
18
19 Set FSO = CreateObject("Scripting.FileSystemObject")
20 Set file = FSO.GetFile(shortcutPath)
21 file.name = bpmf & ".lnk"
```

- WshShortcut simply generates a binary shortcut file as raw bytes that will be written to disk. There are no checks to determine whether the data written is nonsense or will cause problems. In the case of the shortcut's pathname, the file-system APIs reject most nontext characters—and a mangled Unicode filename is highly likely to include some illegal characters.

At this point, we're stuck with WshShortcut' shortcomings. Any chance of WshShortcut being updated to support Unicode vanished when WSH was frozen in 2001. However, as I mentioned previously, there are several workarounds for making Unicode-friendly shortcuts.

Create Once, Copy Many

If you create a shortcut by right-clicking and dragging an object (or by right-clicking in a folder and selecting New Shortcut from the context menu), you'll have no problem including Unicode content in the shortcut's pathname or other properties. If you're creating shortcuts for users on a standardized network (i.e., one that has similar machines, OSs, software, and user setups), you can do this once, put the shortcut on the network (preferably with its read-only attribute set to minimize accidents), then copy it to desktops using a script. Batch files, PowerShell scripts, and scripts using WSH's Scripting .FileSystemObject all support Unicode pathnames, so you don't have any significant restrictions if you use this technique.

Shell Links

The Shell.Application COM object understands both Unicode and Windows shortcuts, which it calls *Shell links*. In spite of this, scripters rarely use Shell links because the syntax is a bit cumbersome and you can't create a shortcut using the scriptable Shell.Application API. However, you can create a shortcut with a tool such as WshShortcut and open it with Shell.Application. You can then use Unicode in any of the shortcut's properties.

I'll show you two ways you can use WshShortcut and Shell.Application to automate the creation of shortcuts that include Unicode. The first technique demonstrates how to use them to create a shortcut with a Unicode filename. The second technique demonstrates how to use them to create a shortcut that has Unicode content in several other properties (e.g., description, target path) as well.

Shell Link Technique 1

If you need to create a shortcut that has a Unicode filename, you can create the shortcut with WshShortcut using an ANSI filename, open it with Shell.Application, then rename the shortcut. Here's a simple demonstration.

Suppose you need to create a shortcut whose filename needs to be in CJK script (core symbols common to the Chinese, Japanese, and Korean languages) on users desktops. You can use a script like UnicodeShortcut1.vbs in Listing 1. Note that this script is displayed in an editor that supports Unicode. The Unicode content can get mangled if the script is opened in an editor that doesn't support Unicode. I use SAPIEN Technologies' free PrimalPad editor. Notepad also supports Unicode.

UnicodeShortcut1.vbs begins by creating the objects it'll use. The script then determines the path to the desktop folder in line 3. In line 4, the script stores the Unicode text you want to eventually use as the filename and description.

Because using a Unicode filename like the one in line 7 (which is commented out) would fail, you need to use an alternate filename that has standard ANSI characters, such as the name *bpmfName* in line 8 (b, p, m, and f are the first four letters in the CJK script system, which is commonly called Bopomofo).

After the shortcut is saved in line 17, it becomes a file. All of the FileSystemObject APIs support Unicode, so you can easily rename the file, as shown in lines 19 through 21. Figure 1 shows the result.

The only thing this technique doesn't do is give you a way to include Unicode content in other shortcut properties. For that, you need to use technique 2.

Shell Link Technique 2

UnicodeShortcut2.vbs in Listing 2 demonstrates the technique for creating shortcuts that include Unicode content in several shortcut properties. I'll walk you through the script so you can see how the technique works.

The script begins by setting four variables:

Figure 1: Shortcut created with UnicodeShortcut1.vbs

Listing 2: UnicodeShortcut2.vbs

UnicodeShortcut2.vbs *

```
1  TargetPath = "C:\Windows\System32\Notepad.exe"
2  Description = "ㄅㄆㄇㄈ"
3  WorkingDirectory = "C:\temp\vm"
4  FinalLinkPath = "C:\tmp\ㄅㄆㄇㄈ.lnk"
5
6  Set WshShell = CreateObject("WScript.Shell")
7  Set ShApp = CreateObject("Shell.Application")
8  Set FSO = CreateObject("Scripting.FileSystemObject")
9
10 ' Create a unique random name in the temp folder.
11 ' Use it to save a shortcut template.
12 tmplnk = WshShell.ExpandEnvironmentStrings("%temp%\") _
13     & FSO.GetTempName() & ".lnk"
14
15 ' Create and save an empty shortcut.
16 WshShell.CreateShortcut(tmplnk).Save()
17
18 ' Open shortcut as Shell Link; the namespace doesn't matter.
19 Set lnk = ShApp.NameSpace(0).ParseName(tmplnk).GetLink
20
21 ' Set properties of shortcut, including new path, and save.
22 lnk.Path = TargetPath
23 lnk.Description = Description
24 lnk.WorkingDirectory = WorkingDirectory
25 lnk.Save(FinalLinkPath)
26
27 ' Force delete the shortcut template.
28 FSO.DeleteFile tmplnk, True
```

- TargetPath, which contains the pathname of the application or file you're creating the shortcut for
- Description, which contains the text that pops up when the mouse pointer hovers over the shortcut
- WorkingDirectory, which contains the working directory for the application or file when it launches
- FinalLinkPath, which contains the location where you want the shortcut to be

Next, the script initializes the objects it'll use (lines 6 through 8) and creates an ANSI path to a location where it can save a shortcut template (lines 12 and 13). This is a template because it has the form of a shortcut, but it is completely empty. The sole reason for creating it is that the Shell .Application object we'll be using for the rest of our work can't create a shortcut file itself; Shell.Application can only modify a pre-existing shortcut.

To create the shortcut template, the script calls the WshShell object's CreateShortcut method with the safe path as its argument. It then immediately calls the returned WshShortcut object's Save method to save it. The script doesn't need to write anything to the shortcut file because WshShortcut doesn't validate content. In fact, it doesn't even give the WshShortcut object a name since it doesn't need one.

The most efficient approach for handling shortcuts that contain Unicode depends on your situation.

In line 19, the script uses the Shell .Application reference in the ShApp variable to get a Shell link object. This line of code starts by connecting to a shell folder's namespace. In this case, the shell folder is the user's desktop folder (indicated by the identifier 0), but it can be any shell folder. The code then calls the shell folder's ParseName method with the pathname to the temporary shortcut as its argument, which provides a connection to that file. The code then retrieves the file as a Shell link by way of the GetLink property.

In lines 22 through 25, the script fills in the shortcut's details and saves the Shell link using the pathname that contains Unicode. Finally, it deletes the shortcut template. Although it takes multiple object references, the entire process can be automated.

The Best Approach for You

The most efficient approach for handling shortcuts that contain Unicode depends on your situation. If you have a standardized network, manually creating a shortcut, then using a script to copy it is probably going to be the simplest solution. This is particularly true if you have only a couple of shortcuts to deploy.

If you don't have a standardized network and you only need to use Unicode in the shortcut's pathname, the first Shell link technique I showed you is probably best. It requires less scripting and customization than the second Shell link technique.

The second Shell link technique is most useful when you need to use Unicode in other shortcut properties, such as its description or working directory. This technique requires the most scripting know-how.

You can download the UnicodeShortcut1.vbs and UnicodeShortcut2.vbs scripts by going to www.windowsitpro.com, entering 125987 in the InstantDoc ID box, clicking Go, then clicking the *Download the Code Here* button. As I mentioned previously, you need to use an editor that supports Unicode (e.g., PrimalPad, Notepad) to open these scripts.

If you're using an editor that doesn't support Unicode, I included alternate versions of these scripts in which the Unicode characters are entered in escaped form. The AlternateUnicodeShortcut1.vbs and AlternateUnicodeShortcut2.vbs scripts work exactly the same way as their counterparts, UnicodeShortcut1.vbs and UnicodeShortcut2.vbs.

InstantDoc ID 125987

Alex K. Angelopoulos (aka@mvps.org) is an IT consultant, an MCSE, and a contributing editor for *Windows IT Pro*. As an avid scripter, he regularly writes about administrative automation using WSH, PowerShell, and related technologies.

PROBLEM:

You run a mix of Windows, UNIX, and Linux systems, and you need to consolidate management and log collection.

SOLUTION:

Use SCOM 2007 R2 with cross-platform support and Audit Collection Services.

WHAT YOU NEED:

SCOM 2007 R2, the downloadable cross-platform management packs, and Audit Collection Services; servers for cross-platform management and collection; and a database server running SQL Server 2008 with SQL Server Reporting Services.

SOLUTION STEPS:

1. Install SCOM 2007 R2
2. Install and configure Audit Collection Services
3. Install cross-platform management support
4. Deploy cross-platform management packs
5. Configure cross-platform Audit Collection Services
6. Tune Audit Collection Services

DIFFICULTY:

Configuring Cross-Platform Support for System Center Operations Manager 2007 R2

Exciting new features let you manage your UNIX and Linux systems alongside your Windows systems

by John Howie

Many enterprises today run a mix of Windows servers and desktops alongside UNIX and Linux systems, with different strategies for managing them. In most cases, the Windows systems are placed into forests and are centrally managed, but the UNIX and Linux systems might or might not be centrally managed. Rarely are Windows systems managed alongside UNIX and Linux systems. For enterprises that do want to manage Windows systems alongside UNIX and Linux systems, there have always been third-party solutions available, but these are often complicated to install and unwieldy to use. These third-party solutions can also be costly and require significant investment in training and staffing to use successfully.

Microsoft has recognized the need to manage UNIX and Linux systems alongside Windows systems, and has over the years provided a means to integrate the systems to provide centralized authentication, authorization, and auditing. With the release of System Center Operations Manager (SCOM) 2007 R2, Microsoft provides support for managing select UNIX and Linux systems through SCOM, as well as extending Audit Collection Services (ACS) to integrate UNIX and Linux event collection, processing, and storage with that for Windows systems. In this article, I describe how to configure and use these exciting new features to manage your UNIX and Linux systems alongside your Windows systems.

Prerequisites

The ability to manage Windows, UNIX, and Linux systems together requires that you have SCOM 2007 R2 deployed in your organization. To integrate event log collection from your Windows systems with your UNIX and Linux systems, you also need to install ACS. In addition, only a subset of common UNIX and Linux systems are supported.

The supported UNIX and Linux systems are AIX 5.3 and 6.1 (Power PC); HP-UX 11iv2 and 11iv3 (PA-RISC and IA64); Red Hat Enterprise Server 4 and 5 (x86 and x64); Solaris 8 and 9 (SPARC) and 10 (SPARC and x86 later than 120012-14); and SUSE Linux Enterprise Server 9 (x86), 10 SP1, and 11 (both x86 and x64). You'll also find that derivatives of enterprise versions of Linux—such as openSUSE—will work, but they're unsupported.

You need to install support for Web Services Management (WS-Man) 1.1 on the Windows servers that host your SCOM servers (which will manage the UNIX and Linux clients). On Windows Server 2008 R2, this is a feature called WinRM IIS Extension that you can add. You also need to install IIS.

I recommend that you install the latest cumulative update for SCOM 2007 R2. You can find the latest update available at the Microsoft Download Center (download.microsoft.com) by searching for the keywords "SCOM cumulative update." The cumulative updates address issues with the use of Server 2008 R2 and SQL Server 2008, and they contain fixes that address many other problems. You need to apply the latest update to your SCOM 2007 R2 Root Management Server(s), any SCOM Gateway servers you might have, and every other SCOM server, as well as all ACS Collectors. You also need to follow instructions for how to update the SQL Server databases that SCOM and ACS use.

You need to install the latest cumulative update for cross-platform support in SCOM 2007 R2. You can also find this update at the Microsoft Download Center by searching for the keywords "cross platform." Unlike the cumulative update for SCOM 2007 R2 itself, there are separate downloads for an SCOM server and an SCOM gateway server.

You need to download the appropriate cross-platform cumulative updates and install them, beginning with your SCOM 2007 R2 Root Management Server (RMS), then your gateway servers, and then every other SCOM server. Read the release notes carefully before applying the cross-platform cumulative update.

Finally, make sure you download the latest cross-platform management pack(s). Currently, there's an installer MSI file and five supporting documents for AIX, HP, Red Hat, Solaris, and SUSE flavors of UNIX and Linux available at the Microsoft Download Center. (Use the search keywords "cross platform.") Review the documents appropriate for the flavors of UNIX and Linux you intend to manage.

The actual management packs are contained in the installer file. Double-click the installer file so that the management packs are extracted and written, by default, to a folder called SCOMCrossPlatformCU2MP, under C:\Program Files\System Center Management Packs. On 64-bit installations of Windows Server, the Program Files (x86) folder is used instead.

Specifying UNIX Accounts

SCOM 2007 R2 uses accounts to monitor and manage UNIX and Linux systems in much the same way as it does Windows systems. SCOM 2007 R2 uses two accounts with UNIX and Linux systems. The first is called a UNIX Action Account and is supposed to be a low-privileged account. The second is called the UNIX Privileged Account and—as the name suggests—is supposed to be a superuser (or root) account.

The majority of UNIX and Linux flavors recognize only two types of users: superusers and ordinary users. Superusers are identified with a user identifier (UID) of 0, such as the user *root*, whereas ordinary users are identified with a UID of any value other than 0. When you're initially configuring SCOM 2007 R2 Action and Privileged accounts, I recommend that you use only superuser accounts for the UNIX Action and Privileged accounts. Once you have SCOM 2007 R2 successfully managing UNIX and Linux systems, you can adjust the credentials associated with the UNIX Action Account.

Now, follow these steps to begin the first phase of your project:

1. To specify credentials for the UNIX Action and Privileged accounts, open the SCOM 2007 R2 Operations Console, select the Administration view, and click Accounts under Run As Configuration.
2. Right-click in the Accounts pane, and select Create Run As Account from the context menu to launch the Create Run As Account Wizard.
3. On the Introduction page (if displayed), select Next.
4. On the General Properties page, select Basic Authentication from the Run As Account drop-down box, type text similar to UNIX/Linux Privileged Account in the *Display name* field, and click Next.
5. On the Credentials page, you need to specify the username and password of the account, then click Next. The account must already exist on your UNIX/Linux systems, or in a centralized directory that they employ (such as Kerberos or NIS+).
6. On the next page, Distribution Security, ensure that the *More secure* option is selected, and click Create. The credentials you just specified will be displayed in the Accounts pane under the category Type: Basic Authentication. Right-click it, and select Properties.
7. In the Properties dialog box, select the Distribution tab and use the Add button to specify the name of every SCOM server that will manage UNIX and Linux systems.

The next step is to associate the credentials you just created with the UNIX Action and Privileged accounts.

1. Click the Profiles node under Run As Configuration in the Administration view, and double-click the UNIX Action Account entry in the Profiles pane.
2. Click Next on the Introduction page (if it appears).
3. On the General Properties page, click Next.
4. On the Run As Accounts page, use the Add button to display the Add a Run As Account dialog box. In the dialog box, select the account you just created from the *Run As account* drop-down list, and make sure the option *All targeted objects* is selected. Close the dialog box and select Save to associate the Run As accounts with the UNIX Action Account.

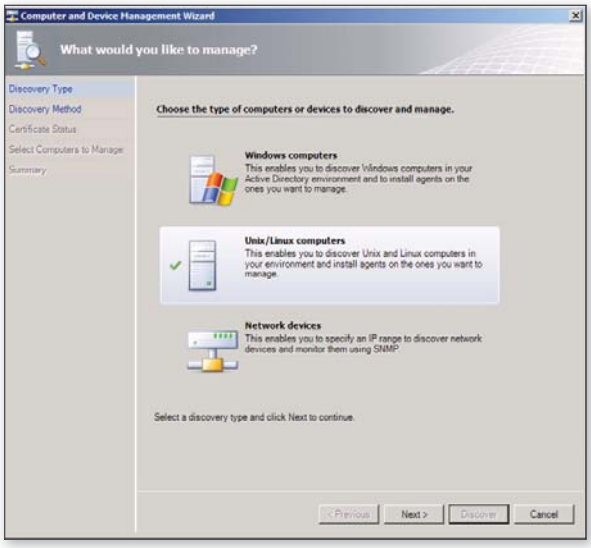

Figure 1: Selecting *Unix/Linux computers* in the Computer and Device Management Wizard

5. When the association is made, you'll be prompted with a warning stating that objects might not be monitored if credentials aren't distributed. As long as you selected the SCOM servers to which the account should be distributed when you created the account, you can ignore this warning.

6. Repeat these steps for the UNIX Privileged Account.

It's important that you remember that—with this configuration—you're using a superuser account for both the UNIX Action and Privileged accounts. After you verify that you can discover and manage your UNIX and Linux systems, you should change the configuration and use a non-privileged account for the UNIX Action Account, if possible.

Deploying Agents with the Wizard

For SCOM 2007 R2 to manage UNIX and Linux systems, they must first be discovered and management agents must be deployed. SCOM 2007 R2 includes a Computer and Device Management Wizard, which you can use to discover UNIX and Linux systems.

1. To find the wizard, open the Administration View in the SCOM 2007 R2 Operations Console, right-click any node, and select Computer and Device Management Wizard. Then, select *Unix/Linux computers*, as you see in Figure 1.

2. Clicking Next will take you to the Discovery Method step, where you can specify how to discover UNIX and Linux systems on your network.

3. Click Add to launch the *Define discovery criteria* dialog box, which Figure 2 shows. You are given three options for discovering UNIX and Linux systems: by IP address, by DNS name, or by IPv4 address range. You also need to specify credentials to scan the network for systems. This is, by default, a superuser account (e.g., the root account), but you can specify an ordinary user account for discovery purposes. If you do, you must clear the *This is a superuser account* option and specify the root password. Clicking OK will save the discovery criteria.

4. You can click Add again to add more discovery criteria. Note that, by default, SSH discovery is disabled. If you leave SSH discovery disabled, the Discovery Agent will only report on UNIX and Linux systems that it finds but won't attempt to install SCOM 2007 R2 cross-platform agents to manage the systems. If you enable SSH discovery, however, SCOM 2007 R2 will attempt to log on to each system it can connect with, using the credentials that you specified. If someone were able to set up a rogue system on your network—within the discovery criteria you specified—that person might be able capture the credentials used in discovery.

5. When you're done adding criteria and credentials, simply click Discover to find UNIX and Linux systems. If the wizard finds systems, and if SSH discovery is enabled, the wizard will attempt to deploy agents to the UNIX and Linux systems. The discovery results will provide details about what systems were found and whether the agents could be deployed.

6. For systems that the agent could be deployed to, select the check box next to the system name and click Next. For systems that were found but weren't fully discoverable, you can click Details to get more information about why the discovery process couldn't be completed. You might find that for systems that weren't discovered initially, rerunning the Computer and Device Management Wizard can yield success on subsequent attempts.

For the systems you selected, the Computer and Device Management Wizard will deploy and install the cross-platform agent for the architecture and install an X.509v3 certificate that the cross-platform agent will use to identify the managed UNIX or Linux system and to secure communications with the SCOM 2007 R2 infrastructure.

Occasionally, a problem can occur with the certificate creation and installation process—typically because of a mismatch in the system's hostname and its DNS name. If a problem with the certificate is reported, you can follow the guidance in the Microsoft article "The Certificate Name Does Not Match the Hostname" (go.microsoft.com/fwlink/?LinkId=148011) to fix the problem, and rerun the Computer

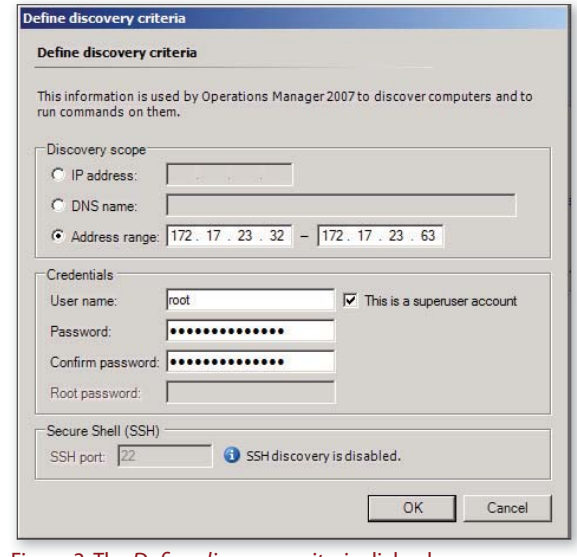

Figure 2: The *Define discovery criteria* dialog box

and Device Management Wizard using the process that I described earlier for the affected systems.

Manually Deploying Agents

If you can't use the Computer and Device Management Wizard to install agents to your UNIX and Linux systems, you can manually install them. You'll find the agents for each supported platform in the folder C:\Program Files\System Center Operations Manager 2007\AgentManagement\UnixAgents. Copy these agents to an FTP server or website so that you can download them to your UNIX and Linux systems. You can also add these to baseline images if you use them in your organization. If you applied the latest cumulative update for cross-platform support, you'll find that there are different versions of the agents in this folder. You should always use the latest agents. For information about how to install the agents for each supported platform, see the Microsoft article "Manually Installing Cross Platform Agents" (technet.microsoft.com/en-us/library/dd789016.aspx).

Once you've manually installed the agents onto your UNIX and Linux systems, you'll need to rerun the Computer and Device Management Wizard using the process I described. The wizard will find the systems with agents manually installed and ask whether you want to issue new X.509v3 certificates to them. Select the systems you want to install certificates to, and click Sign. Once the certificate(s) have been issued, the Computer and Device Management Wizard continues, and you need to select the system(s) you want to add to the pool of managed UNIX and Linux servers in SCOM 2007 R2—in a process similar to the automatic discovery of UNIX and Linux servers.

Monitoring UNIX and Linux Systems

With agents installed, you can begin to monitor your UNIX and Linux systems from the SCOM 2007 R2 Operations Console. Simply select the Monitoring view, as you see in Figure 3, then click the Unix/Linux Servers node. When you select a system in the Unix/Linux Servers pane with a cross-platform agent installed, you'll see a summary of the information SCOM 2007 R2 has about the system listed in the Detail view. You can also use the Health Explorer to analyze the system and core processes, such as Cron, SSH, and Syslog. The information available in the Health Explorer varies depending on the target system. You can also put a system with the cross-platform agent installed into Maintenance Mode, much like a regular Windows server. And in the State view, with the cross-platform agent installed, you can also run tasks. There are three tasks available: Memory Information, Run VMStat, and Top 10 CPU Processes.

Also in the Monitoring view, you can diagram and view other information about your UNIX and Linux servers, and you can configure performance monitoring. Expand the Unix/Linux Servers folder, expand the OS folder beneath, and select the appropriate nodes. The nodes available and the data returned will vary by OS type and depend on support in the cross-platform agent and appropriate management packs installed.

As more cumulative updates are released for cross-platform support, or as third-party management packs are released, the ability of SCOM 2007 R2 to manage UNIX and Linux systems will increase. However, if you download new agents or management packs, you need to rerun the Computer and Device Management Wizard to deploy these updates or to sign certificates for agents you manually deploy.

Configuring ACS for Cross-Platform Support

To configure ACS for cross-platform support, you need to perform several steps on both your SCOM infrastructure and your UNIX and Linux servers. You can turn on cross-platform support for ACS only if you already have ACS installed and configured—including ACS Reports.

1. You need to download the latest ACS cross-platform support software by visiting the Microsoft Download Center and searching for "ACS cross platform." You need to download both the Cross Platform Audit Collection Services software, which consists of 32-bit and 64-bit installers and supporting documentation, and the Cross Platform Audit Collection Services Management Packs.

2. You need a server that will act as a collector of audit events from your UNIX and Linux systems and forward them to ACS. This server must be configured as an ACS Collector. You might want to consider creating dedicated SCOM Management Servers for your UNIX and Linux hosts, and make them ACS Collectors, too. On this server, you need to install the Cross Platform Audit Collection Services MSI file that you downloaded. Double-click the MSI in Windows Explorer to begin installation, and accept the license agreement.

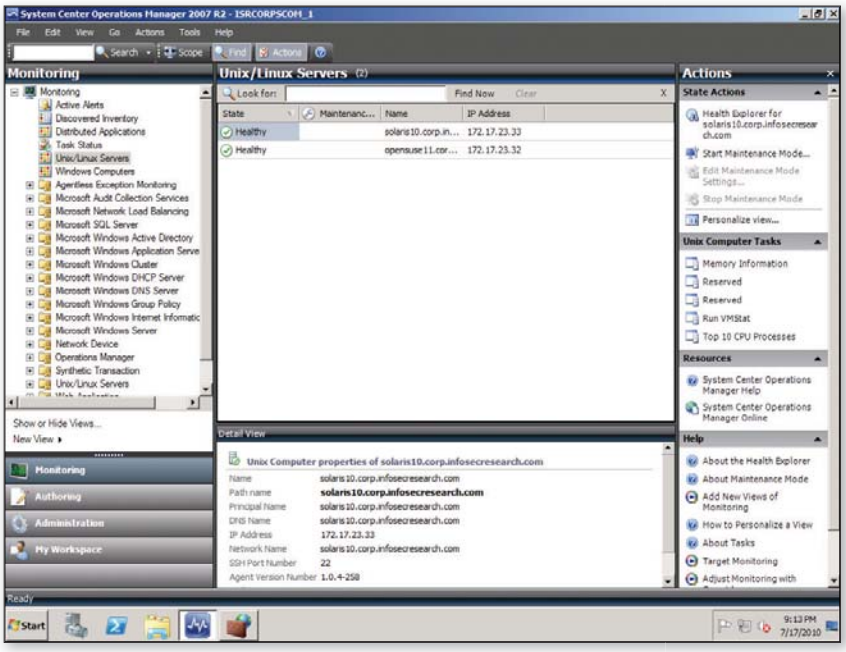

Figure 3: Monitoring UNIX/Linux servers

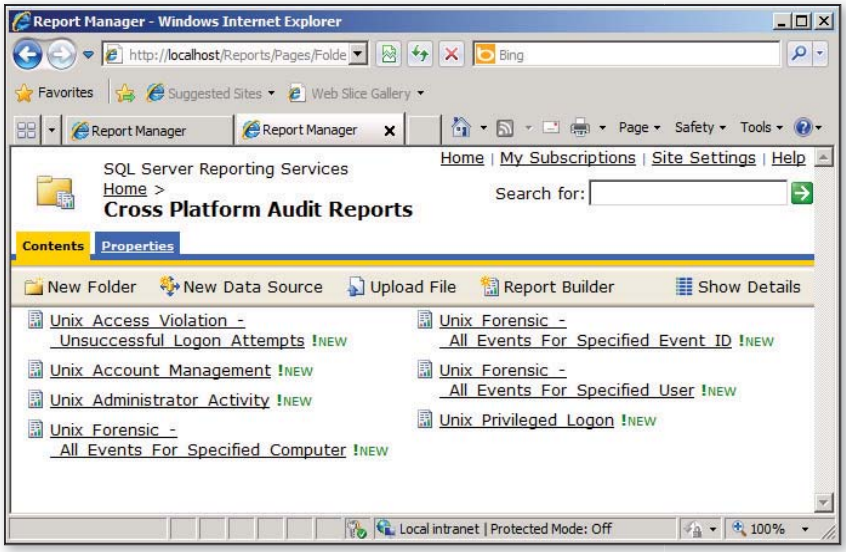

Figure 4: Cross Platform Audit Reports

3. In the Audit Data Time Zone step, select either the current local time or the Coordinated Universal Time (UTC). You should pick the same option you used on your ACS Collectors, which is probably UTC.

4. The last page of the installation wizard will caution you that cross-platform support for ACS can generate a lot of events, which might stress your ACS database. The actual volume of events generated will depend on many factors, including the number of UNIX and Linux servers you have and how you configure each. You can tune performance if necessary by using the standard ACS management tools. If you have many UNIX and Linux servers, you might need to install additional cross-platform collectors, depending on performance and fault tolerance. Finally, make sure that the Group Policy setting *Audit object access* is set to Success and Failure, if you're using it.

5. Next, you need to import the ACS cross-platform management packs you also downloaded. Run the installer to extract the management packs to the folder \System Center Management Packs\Operations Manager 2007 R2 Cross Platform ACS MP—under C:\Program Files or C:\Program Files (x86)—and import them into SCOM using the Import Management Packs wizard. If your SCOM 2007 R2 system is connected to the Internet, you can also download the management packs as you would any other management pack.

6. Now, it's time to install the cross-platform audit reports to your ACS Reporting Server. On the dedicated ACS cross-platform collection server, log on as a user who has administrative privileges to the SQL Server Reporting Services (SSRS) server used by ACS, open a command prompt, navigate to the folder C:\Program Files\System Center Operations Manager Cross Platform ACS, and run the following command:

```
.\UploadCrossPlatformAuditReports
  .cmd "<Audit DB Server>\<DB
  Instance>" "http://<Report Server>/
  ReportServer[$<SRS Instance>]"
  "C:\Program Files\System Center
  Operations Manager Cross Platform
  ACS"
```

7. Replace *<Audit DB Server>* with the name of your ACS database, *<DB Instance>* with the instance name of SQL Server (the default is MSSQLSERVER), *<Report Server>* with the name of your SSRS server, and append *$<SRS Instance>* where *<SRS Instance>* is the name of the SSRS instance if you're using anything other than the default instance. After installation, the reports added are written to the folder Cross Platform Audit Reports, as Figure 4 shows.

8. At this point, you need to enable cross-platform support for ACS. To do so, you need to modify management packs. Best practice will have you modify copies of management packs that you create for modification. To modify management packs for ACS, select the Authoring view in the Operations console and expand the Authoring node, then the Management Pack Objects node, and select Object Discoveries.

9. In the Object Discoveries pane, search for ACS. A list of ACS endpoints for the various supported UNIX and Linux systems appears, but you care only about the entry *Discovered Type: ACS Endpoint*. Right-click *Discover Unix/Linux ACS Endpoint* underneath the entry, and select Overrides, then *Override the Object Discovery*, then *For all objects of class: Unix Computer*.

10. In the Override Properties dialog box, select the Override check box, click Apply, then click OK.

For most enterprise deployments of Linux, there's no need to configure the servers or Syslog, and security events of interest will start to flow into ACS; you can view SSRS on your ACS report server.

If you have nonstandard Syslog configurations and are using Rsyslogd or Solaris or AIX, you need to configure Syslog to write security-related events to /var/log/messages (for Linux-based systems). For Solaris and AIX systems, you need to follow the guidance available in the Microsoft article "Configure Syslog and Rules for Audit Events" (technet.microsoft.com/en-us/library/ee909515.aspx).

Centralized Monitoring

The steps I described to get cross-platform support up and running for SCOM 2007 R2 aren't easy, and you might find that it takes some experimentation to get everything working correctly. That's especially true for ACS. However, the return on the time invested in getting integration working will pay off as you find that you can monitor your Windows, UNIX, and Linux systems from one place.

InstantDoc ID 125988

John Howie
(jhowie@microsoft.com) is a senior director in the Online Services Security & Compliance team at Microsoft, where he manages cloud security.

EDRM and SharePoint

Designing and Building a Compliant Platform

Create a trusted central document repository

by Ron Charity

The amount of information generated by organizations is growing at a phenomenal rate. With factors such as increased legislation and penalties, many organizations want to leverage their information assets, both physical and digital, to their full advantage by increasing speed of access, improving classification schemes, and ensuring information quality and trustworthiness. To do so, organizations are adopting toolsets such as Electronic Document and Records Management (EDRM) applications, Microsoft SharePoint, and other document management tools for creating a trusted central repository, which is a key goal of records managers.

There are many roadblocks to attaining that goal, such as executive sponsorship and enforcement, organizational alignment, confusion over toolsets, lack of a consistent and user-relevant classification scheme, poor usability, and low user adoption. Most organizations have islands of EDRM systems in various states of deployment that are generally not delivering the value they promised. That's why having a good planning and deployment strategy in place is crucial.

Planning Your Records Management Solution

Planning a records management solution is a daunting task. To be successful, you need a multidiscipline team because designing the system requires multiple skills and diverse knowledge (e.g., communications, governance, information architecture, records management, data migration, security, EDRM system architecture). A good place to start the planning process is to read ISO 15489 and the DIRKS methodology. (See the sidebar "EDRM Related Resources.") Key points that you need to consider when planning your EDRM rollout include the following:

- Governance—You need an executive sponsor to drive your records management program across the organization.
- Team—You need to work with a multidiscipline team as discussed earlier.
- Plan—You must have a plan that will deliver the expected value for the organization.
- User adoption program—You must create a program that encourages and measures user uptake, provides education and mentoring, and enforces best practices (performance reviews).
- Data policy—You must define, communicate, and enforce the data policy for records. People must understand their roles and responsibilities.
- Inventory of data repositories—You need to conduct an inventory of what you have, the state it's in (age, usefulness), and location. The inventory might include file shares, public folders, microfiche, third-party records storage warehouse, and departmental EDRM systems.

Records Center 2007 Overview

Although Microsoft has released SharePoint 2010, there is merit in covering SharePoint 2007 technologies because many organizations won't upgrade for another year or so. SharePoint 2007 has a site template called Records Center. It has functionality such as Holds, which lets you place holds on

retention schedules, and Auditing, which logs events and operations against the record. It also contains some specialized libraries and content type routing configuration that allows you to route records by content type.

For example, you could create a content type called Contracts and a library for the content type. You would then create a routing rule in Records Center that would route the content type to the Contracts library. You could then apply Information Management policy (i.e., how to handle records) to the library. The final step in the process is to configure the farm to be aware of the Records Center by adding the link to the Records Center web service (e.g., *server_name*/_vti_bin/*officialfile*.asmx). To do so, select External Service Connections and enter the URL to your records center web service in the *Connect to a Records Repository* field. Upon completion, the option to route records within document libraries will appear in the item's edit/context menu as Figure 1 shows. Note that any work you do should be done in a sandbox environment, which is usually a virtual machine (VM) running on your laptop.

Create a Records Center

This section outlines the basic process for creating a Records Center in Microsoft Office SharePoint Server (MOSS) 2007. For more detailed information, see the Microsoft TechNet article "Create a site collection (Office SharePoint Server)" at technet.microsoft.com/en-us/library/cc263094(office.12).aspx. To create a Records Center, follow these steps:

1. Log on as Administrator and open Central Administration. Select Application Management.
2. Enter a name for the application, SQL Server credentials, and authentication method that suits your environment. Using Quotas is also a best practice but isn't a requirement for a sandbox environment. Note that isolating Records Center with its own site and application will help simplify operations.
3. Under Site Collection, select *Create a Site Collection*. Enter a name, description, and URL, then select the Records Center template. (Choose a URL and port number that fits with your environment. Also, associate your Site Collection with the application you created in the prior step.)
4. Select Application Management and click External Service Connections. Enter the URL to your Records Center web service in the *Connect to a Records Repository* field (e.g., *server_name*/_vti_bin/*officialfile*.asmx). Note that in SharePoint 2007 you can specify only one connection per farm.
5. Open your Records Center site in the browser and click Site Settings. Create a content type called Contracts, and accept the default settings. Click Site Settings, select Create, and create a document library called Contracts.
6. Configure the Contracts document library to accept multiple content types.
7. Create a routing rule called Contracts that directs the Contracts content type to the Contracts document library by specifying the name of the library as the destination.

You now have a functional Records Center site. If desired, you can add the Records Site to your Portals navigation by editing the Portals navigation menu and adding a tab for the new Records Center site.

Records Center 2010 Overview

In SharePoint 2010, the Records Center has a slick new UI and some additional features that improve usability and management. For those who evaluated SharePoint 2007 and held off, SharePoint 2010 offers compelling new features and enhancements. Here is an overview of some of the new functionality.

In-place records management. A new industry trend is the idea of in-place records management. Specifically, instead of a central repository of documents that requires a routing service or manual movement of files, records are managed where they are. The documents stay in the current location and are classified as business records. This approach avoids the laborious migrations associated with EDRM projects.

Multistage retention. Retention policies can have multiple stages, allowing you to specify the entire document lifecycle as one policy (e.g., review Contracts every year, and delete after seven years).

Folders. Another interesting note is the new role of folders in SharePoint 2010. In SharePoint 2007, folders were the recommended option for grouping files by security access. With SharePoint 2010, folders can act as true parents to any child objects below them. Now, metadata is set at the folder level so that the child objects can inherit that information.

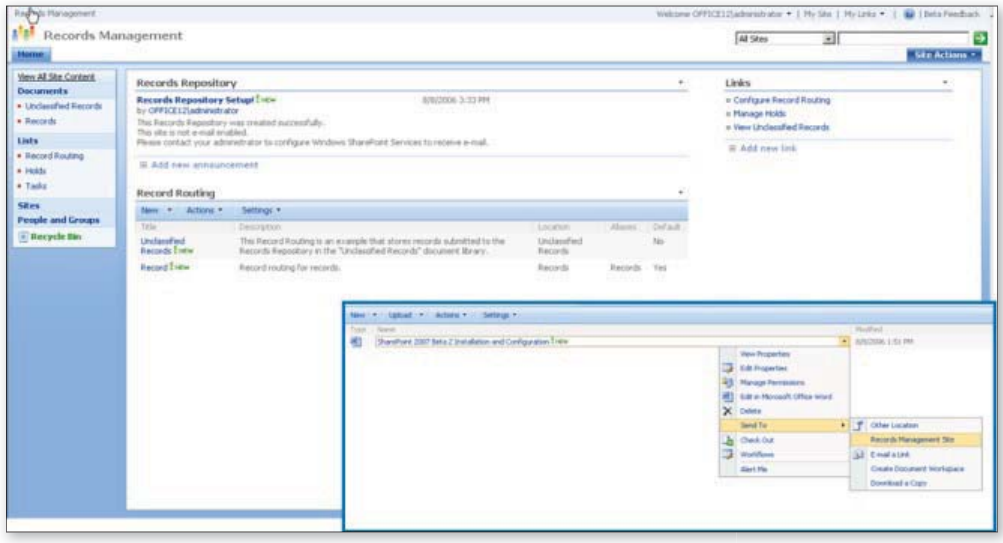

Figure 1: Records Center

Document sets. Document sets enable users to collaborate on related documents without having to create a new document library or site. Their purpose is to help organize documents and enable collaboration on documents tagged with similar metadata.

Persistent document IDs. The document ID provides absolute reference to objects, regardless of filename changes or document moves.

Content organizer. The routing rules from SharePoint 2007 have been replaced by the content organizer, a new SharePoint feature available in all document libraries. The content organizer routes documents to the correct folder based on content types and any other metadata that you require.

Compliance details. This is a new feature that adds a Compliance Details option to the context menu. It allows users to check out the relevant settings that have been applied to a specific business record. The feature also enables administrators to make sure that specific documents are inheriting the correct policies and retention settings.

Submit a record. This is a basic feature that adds some simple usability to the site.

For more information, read the article "8 Reasons SharePoint 2010 Looks Like a True ECM System" (aiim.typepad.com/aiim_blog/2010/01/8-reasons-sharepoint-2010-looks-like-a-true-ecm-system.html).

Creating a Records Center in SharePoint 2010

The following steps provide the basic outline for creating a Records Center in SharePoint 2010. To find more in-depth information, see the TechNet article "Create a site collection (Office SharePoint Server)" at technet.microsoft.com/en-us/library/cc263094(office.12).aspx.

1. Log on as Administrator and open Central Administration. Select Application Management.
2. Open the Create Site Collection page. In the Web Application section, if the web application in which you want to create the site collection isn't selected, on the Web Application menu click Change Web Application, then click the web application in which you want to create the site collection. This process is detailed in the TechNet article "Create a site collection (SharePoint Server 2010)" at technet.microsoft.com/en-us/library/cc263094.aspx.
3. As with SharePoint 2007, create a web application for Records Center. Chose a name, SQL Server credentials, and an authentication method that suits your environment.
4. On the Create Site Collection page, in the Web Application section, if the web application in which you want to create the site collection isn't selected, on the Web Application menu, click Change Web Application, and click the web application in which you want to create the site collection.
5. In the Title and Description section, type the title and description for the site collection. In the Web Site Address section, select the path to use for your URL (e.g., a wildcard inclusion path such as /sites/, or the root directory). In the Template Selection section select Records Center from the list. Click OK.

Note that SharePoint 2010 enables administrators to create multiple connections to a Records Center per farm, and to specify more detailed parameters about the behavior of those connections. Another enhancement is the ability to move an item to the Records Center, or to move an item but leave a link in the original location.

Next you need to create content organization rules. You add rules by following these steps:

1. Select Site Settings, Site Administration, Content Organizer Rules. Then select New to open the New Rule page.
2. Create a rule called Contacts. Page down and select the target location. Browse to a library in your site. Make sure that whatever content type you choose is activated in the target library, or you'll receive an error that the destination doesn't have the content enabled.

It's important to note that the source and destination libraries (i.e., document libraries configured for in-place records management) require the same content types. For example, if you have a Contracts Form content type in your source library, you also need one in your destination library.

Figure 2: In Place Records Management feature

In-Place Records Management

To take advantage of SharePoint's In Place Records Management feature, you must first activate the feature for the site collection, as Figure 2 shows. If you have a site collection configured in your sandbox environment, go to Site Actions and activate the feature. Otherwise, create a site collection (e.g., Team Collaboration), then activate it.

If you want items within the libraries to be automatically flagged as records (e.g., a Contracts library), choose the library in your site that you want to enable to support in-place records management, and click Library Settings. Enable Automatic Declaration. Now, all documents added to the library are automatically flagged as records. Note that after you activate the feature and configure the library, a new button, Declare Records, appears in the Libraries ribbon. You can read more in the TechNet articles "Records management overview (SharePoint Server 2010)" (technet.microsoft.com/en-us/library/cc261982.aspx) and "Designing for in-place records management" (technet.microsoft.com/en-us/library/ff363732.aspx).

Third-Party EDRM

Butler Group released a review of the top records management vendors and their products in late 2008. EMC Documentum, IBM FileNet, and OpenText were the notable top players. HP TRIM is classified as Consider, and SharePoint 2007 is classified as Explore.

An important pattern to note is the rise of SharePoint. It's disrupting the EDRM market, as well as some others, such as Web 2.0. With its high cohesion with Office, its ease of use (generally a high rate of adoption), and its low cost from a licensing

EDRM Related Resources

Are you in the midst of planning your EDRM project or looking to improve your deployment? Or perhaps you want to improve operational efficiency? Here are some related resources that will help you on your way.

- **AIIM:** An industry consortium of vendors and professionals focused on establishing standards and professional practice in the field of information management (www.aiim.org).
- **ARMA:** An industry consortium focused on establishing standards and professional practices in the field of information management (www.arma.org).
- **ISO 15489:** A standard based on the Australian Standard AS 4390 (www.iso.org/iso/catalogue_detail.htm?csnumber=31908).
- **DIRKS:** A methodology developed by the Australian Government for planning, deploying, and operating EDRM systems. This should be your go-to guide (www.naa.gov.au/records-management/systems/dirks/index.aspx).
- **Microsoft:** Material for planning your deployment and its operation (office.microsoft.com/en-us/sharepointserver/HA101735961033.aspx).

perspective, SharePoint is a good option for more and more companies. According to Butler Group, "SharePoint is designed to help organizations manage content while enabling information workers to communicate and collaborate electronically." EDRM vendors have made attempts to integrate with SharePoint using simple Web Parts that expose the file plan and provide varying levels of search integration. Typical EDRM Web Parts you might see include the following.

Document library. This Web Part enables site administrators to expose the EDRM file plan as a document library. You need to specify the area (Container/Class) of the file plan to expose it. This Web Part can make the provisioning process more difficult because of the skills required to configure it when provisioning sites; therefore, adopt with caution.

Edit/context menu. The edit/context menu Web Part will include the Move Document Feature option. This Web Part adds a feature to the edit/context menu that enables users to select a document to be moved to the EDRM file plan. It generally exposes an ASPX page that requests record type and required metadata before sending to the EDRM.

Search protocol handler. This Web Part is installed on the index server and enables the server to crawl third-party applications such as EDRM. For SharePoint to display combined search results from a content source, it must understand how to access, crawl, and return results. This is the job of the handler. I've found handlers to perform slowly and sometimes hang the indexer. Adopt with caution.

Federated search. If vendors don't have the time or expertise to develop a handler, the federated search Web Part is available. SharePoint search passes the Web Part, receives the search criteria from SharePoint search, processes the query, and displays it as a separate results set. I've found these Web Parts easy to install, and they perform well.

I've worked with a few EDRM products' Web Part integration and wasn't impressed. They don't provide true integration, but instead are more or less a simple gateway to the EDRM via Web Parts. They don't provide the ability to take advantage of features such as alerts, workflows, and views. Indexing performance is generally slow and doesn't scale well. There are often stability concerns and provisioning is problematic. Given its almost viral growth rate, SharePoint will continue to disrupt the EDRM market as it did the Portal market.

To the Future

To understand how communication has changed over the past century, and especially in the past 20 years, think about the communication media that past generations used: letters, phone, fax, and of course face-to-face communication. With these forms of communication came a simpler slow-paced work environment than we experience today. Today, we use a variety of devices and applications such as smartphones, email, IM, cameras, Facebook, LinkedIn, VoIP, web conferencing, Twitter, and corporate systems such as email and voicemail.

EDRM is no longer the simple capture of physical and electronic assets such as mail and electronic documents. The scope of EDRM spans many applications and communication media and is far reaching organizationally. Some organizations realize this and have locked down corporate access to Internet applications. But what are the employees communicating from home? Are they leaking company secrets? Are there inappropriate communications about the organization or its people? Are recruiters chasing your staff on recruitment and networking sites?

Tools such as SharePoint and technologies associated with Web 2.0 have changed the EDRM market. SharePoint 2010 is applying pressure to EDRM systems to step up their game regarding in-place records management. Records managers and IT managers are thinking of ways to capture communication in applications such as Facebook and Twitter. The price point and usability of SharePoint is pushing EDRM tools out and changing the way companies think of and manage information assets. I predict that standard EDRM will become a backend to software such as SharePoint. EDRM vendors will improve their integration by creating deep linkages to applications such as SharePoint and Web 2.0 applications that are seamless to the average user. Users won't even know their documents and communications are being captured as records.

InstantDoc ID 125915

Ron Charity
(ron.j.charity@sympatico.ca) is a SharePoint product manager with a leading consulting firm. He has worked with SharePoint since 2001 and focuses on governance, information architecture, technical architecture, and operations.

- Outlook
- Security
- Cloud Computing
- Patch Management

Digital Signatures now Available as SaaS

Secured Signing is now offering PKI-based secured signing transactions using ARX's **CoSign** digital signatures. With this new partnership, customers will be able to sign and verify signatures in a non-proprietary format. Users will receive an individual signing key that is protected by the CoSign device, giving each user exclusive access to his/her signing key. According to the vendor, this is a value-add from other digital signature services where users are not given exclusive access to their keys. For a free trial of the solution, visit www.securedsigning.com.

EminentWare Releases Patch Management for WSUS or SCCM

EminentWare has released the **EminentWare Extension Pack,** which brings patch management features to Windows Server Update Services (WSUS) or System Center Configuration Manager (SCCM). Capabilities of the product include: patch third-party apps using WSUS or SCCM, receive pre-built third-party patches, determine which patches have been approved and need to be installed on which machines, deploy patches with scheduling options and approval management, generate reports, use Wake-on-LAN to ensure successful installation, perform discovery and inventory system details, and manage enterprise-wide desktop configurations. To learn more, visit www.eminentware.com.

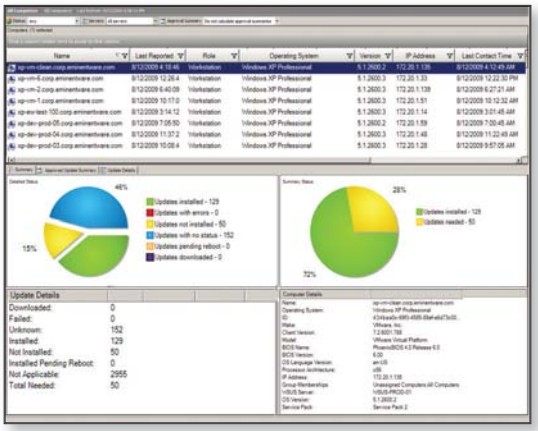

PRODUCT SPOTLIGHT

Messageware OWA Suite 2010 for Exchange Server 2010 SP1

Messageware has announced the general availability of **Messageware OWA Suite 2010** for Exchange Server 2010 Service Pack 1 (SP1). OWA Suite offers add-ons to enhance Outlook Web Access, the web version of Outlook, to create a full-featured Outlook experience to end users.

"Many of our current customers and prospects have been eagerly anticipating the release of the first Service Pack for Exchange 2010 so they can upgrade and deploy Microsoft and Messageware together," said Mark Rotman, president and CEO of Messageware. "In the coming months, we will be working with numerous enterprises throughout the world to realize the additional features and benefits in this latest version."

Aside from bringing existing Messageware OWA Suite functionality to Exchange 2010, there are a few new features in this latest version. One such addition is CalendarPack 2010, which adds new functionality to calendaring in OWA to better match the features in desktop Outlook. Examples of new features include holiday file support, the ability to manage delegate rights from within OWA, and advanced calendar search.

"Microsoft Exchange Server 2010 helps users be more productive by delivering 'anywhere access' to business communications through features like Outlook Web App," said Ian Hameroff, Microsoft's group product manager for Exchange partner marketing. "By working closely with Microsoft and the Exchange 2010 SP1 Beta, Messageware has built on the new features and functionality of SP1, offering Exchange customers the benefits of additional functionality and security for Outlook Web App."

To learn more, visit www.messageware.com.

Djigzo And Comodo Partner For Secure Email Certificate Solution

Djigzo has teamed with Comodo to provide an easy-to-use email encryption and authentication solution. Djigzo's email encryption and authentication gateway will interface directly with Comodo's managed Enterprise Public Key Infrastructure (EPKI). According to the vendors, large organizations that need personal certificates for all employees can now automatically obtain the required certificates without having to use complex procedures. Djigzo's email encryption gateway is open source and available for free at www.djigzo.com. To learn more about Comodo's security solutions, visit www.comodo.com.

EXLADE Updates Cryptic Disk

EXLADE has released **Cryptic Disk 3.0,** the latest version of its disk encryption software. Cryptic Disk lets users create virtual encrypted disks, or encrypt hard disks, drive partitions, USB drives, and memory cards. New features in version 3.0 include: an improved data encryption process, pushing encryption key sizes from 256 bits to 2,944 bits; support for setting automatic actions upon mounting and dismounting an encrypted disk; and a redesigned interface. The product can also create encrypted disks within existing disks, with up to three layers of nesting. To learn more, visit www.exlade.com.

Recover Damaged Outlook Files

DataNumen has released **Advanced Outlook Repair 3.2,** a product that

NEW & IMPROVED

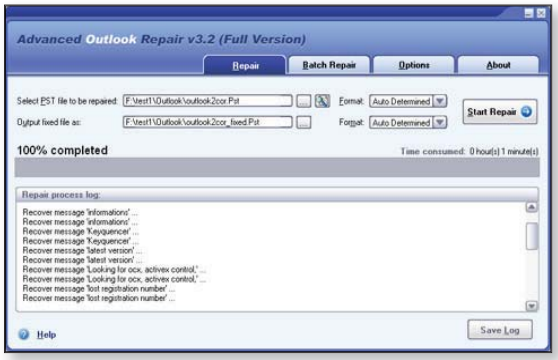

recovers corrupt or damaged Outlook files. According to the vendor, **Advanced Outlook Repair 3.2** is able to recover corrupt or damaged PST files with an average recovery rate of 95.7 percent. In addition to a more robust repair engine, version 3.2 offers improved performance. Advanced Outlook Repair can also recover attachments, embedded objects, folders, appointments, meetings requests, contacts, password-protected emails, and more. To learn more, visit www.datanumen.com/aor.

Enhance Outlook Search with Lookeen

Axonic introduces a product called **Lookeen** to enhance Outlook search. While Outlook's search is already very robust, Lookeen offers the following features: search external PST files and archives; search for attachments, appointments, tasks, notes, and contacts; integrated desktop search; and central indexing for enterprises. To learn more or download a 14-day trial, visit www.lookeen.net.

ElcomSoft Updates Advanced Office Password Recovery

ElcomSoft has announced an update to **Advanced Office Password Recovery** (AOPR), a tool that removes document restrictions and recovers passwords to Microsoft Office documents. AOPR has been enhanced to work with the strengthened Office passwords found in Office 2010 and 2007. According to the vendor, the product can crash up to 20,000 passwords per second in Office 2010 and 40,000 per second in Office 2007. To learn more, visit www.office.elcomsoft.com.

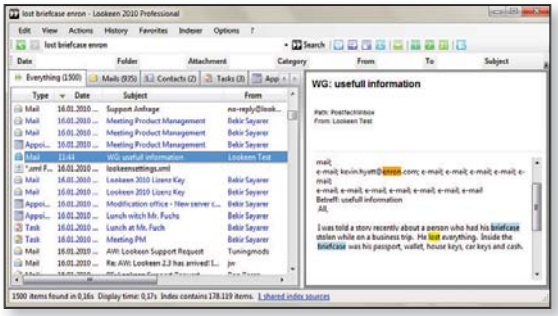

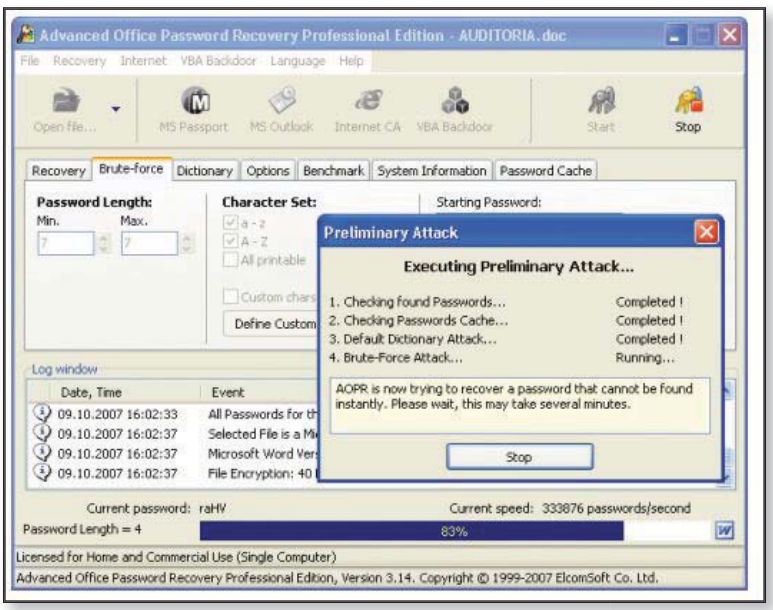

Paul's Picks

www.winsupersite.com

SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

Amazon Kindle 3

PROS: Correctly priced at last; superior screen; smaller, lighter; amazing battery life
CONS: Still no color support
RATING: ◆◆◆◆◆

RECOMMENDATION: Amazon's Kindle 3 is the superior e-book reader solution, with the best screen on the market, offering a 50 percent improvement in contrast that's both immediately noticeable and desirable. And since it uses e-ink technology, the Kindle can be used in direct sunlight, unlike Apple's lackluster iPad, whose screen is ultra-reflective in any light. Performance is up, the device is smaller, lighter, and better looking than its predecessor, and the onscreen display is further improved with better fonts and a nicer layout for periodicals. Yes, the same tired complaints can still be leveled at the Kindle—it has no color output, for example, and no backlighting. No matter: If you care about reading, the Kindle is the obvious choice. It's now available in two versions, a Wi-Fi-only model for $140, and a 3G-equipped model—with lifetime wireless access to Amazon's online store—for $190. These are significantly cheaper than the cost of the Kindle 2 a year ago, which retailed for $260.
CONTACT: Amazon • www.amazon.com
DISCUSSION: www.winsupersite.com/mobile/kindle3.asp

Apple iTunes 10

PROS: Slightly faster than previous versions; faster device sync
CONS: Still a performance dog; lackluster feature set; unsophisticated media management
RATING: ◆◇◇◇◇

RECOMMENDATION: Apple makes tremendous hardware, but its software efforts, especially on Windows, have moved beyond suspect into the ludicrous. Yes, iTunes 10 is a bit faster than its predecessor, and yes, it does offer slightly better device sync with iPods, iPhones, and iPads. But the software is a miserable performer overall and needs an overhaul. And unlike the past few revisions, iTunes 10 doesn't include any major new features. There's a terrible social networking service called Ping that's bundled into the application, and the new default view for music content is still far too text based. Worse, Apple has wiped all color from the application, with new grayscale icons and onscreen elements, creating a washed-out look. For an application that's all about digital media content, iTunes 10 is curiously lacking in emotion or excitement. It's just a terrible mess.
CONTACT: Apple • www.apple.com
DISCUSSION: www.winsupersite.com/digitalmedia/itunes10.asp

InstantDoc ID 126040

November 2010

The Essential Guide to Client Backup for Windows 7

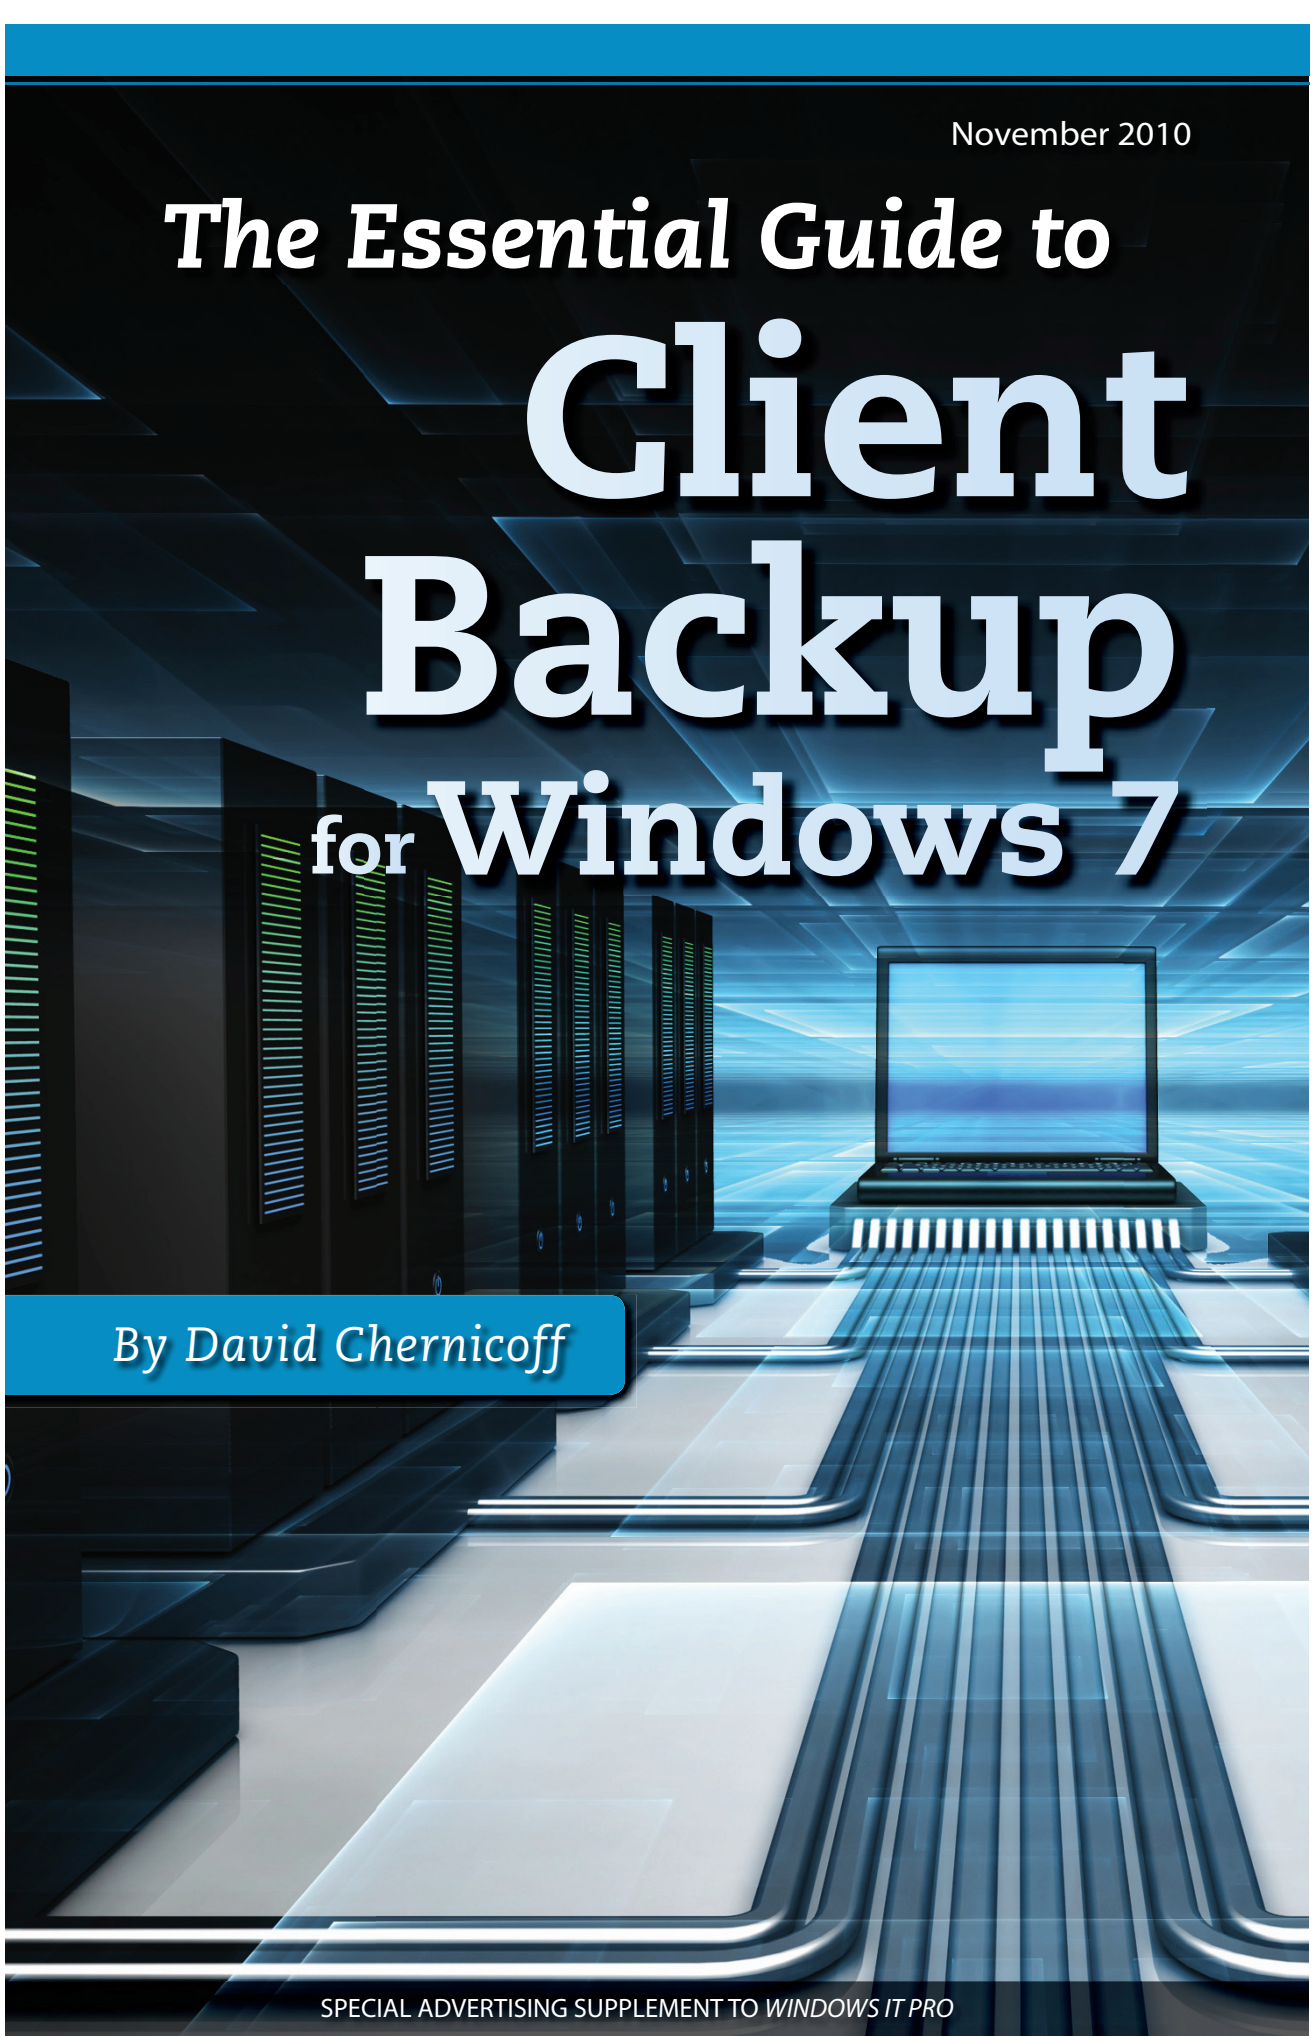

By David Chernicoff

SPECIAL ADVERTISING SUPPLEMENT TO *WINDOWS IT PRO*

There was a time when the idea of client backup meant little more than finding a way to provide data protection for the small hard drives installed on desktop computers in your enterprise. But those days are long gone: The idea of backing up clients has gone from making sure there was a tape backup somewhere of the data on the client computers to providing near real-time backup for every client in the enterprise, be it a desktop or mobile user, regardless of desktop operating system.

Many things complicate the task of providing client backup. First, the amount of data that can be stored on a client computer has increased exponentially. In a day when hard disks of only a few hundred gigabytes is consider small, the average client computer has storage that rivals that found on a server a decade ago. And the price of these drives is so low that it isn't even cost effective for vendors to produce smaller drives that might better suit the needs of corporate desktops. Notebook drives are well over 100 GB, and even netbooks come standard with 160 GB drives.

What this leads to is an immense amount of data being stored on client systems throughout the enterprise. And while it's possible, in some enterprises, to back up only protected user directories to network storage, the proliferation of remote offices and mobile and casually connected users means that there will always be an ongoing need for fast, effective, and efficient client backup. And a backup solution that fits in with the overall corporate business continuity plan. With the widespread adoption of Windows 7, IT has the perfect opportunity to revisit and evaluate their client backup technologies.

Understanding the Problems

The major problem of corporate backup is the presumption that it's being properly done and actually working well. That may sound overly simplistic, but it is fundamentally the issue. Despite years of advice, people still don't think about backup until something goes wrong and they need to do a restore. What this often translates into is including desktop users in the backup process, something that rarely makes for a reliable backup.

Reliable backup and restore for the desktop user

Client backup needs to be completely automated once installed and configured by the IT department. The user should have little to no control over the backup process, nor should the process have a negative effect on the user experience. It can't slow down the computer noticeably or prevent user access to files when needed. And it still needs to reliably back up any changes made to files being protected without user intervention.

But that's not to say the user shouldn't be able to interact with the backup application. Ideally, there is a user-friendly mechanism, be it via application or web browser, which can be configured by IT to allow a user to recover lost files. The goal is to make the data protection process invisible to the user and to give that same user some feeling of control over her data.

Simplifying backup for the mobile user

Mobile clients present their own special needs. In too many organizations, mobile backup is handled directly by the user and is often little more than copying files on the mobile computer to external media of some sort. There is little control over the process and rarely any monitoring that assures that, if necessary, critical data can be recovered. Given the additional data security issues that ad hoc file copying can present, it's clear that a comprehensive client backup strategy needs to be able to deal with the computer that has limited connectivity back to the corporate network.

Assuring a Smooth OS migration

With the widespread acceptance of Windows 7 in the corporate environment, many organizations are looking at their first major operating system migration in many years. Given the nature of the Windows 7 upgrade, this is the perfect time for IT to re-evaluate its enterprise backup infrastructure. Clearly, the need to fully support the new operating system is paramount, but many other factors need to be considered, starting with the ability to ensure that critical client data has been backed up, and is restorable, before and after the migration to the new operating system.

But even if a full operating system migration is set for the near term, many normal activities, such as patch and hotfix application, to both operating systems and applications, should be recoverable with a good backup tool. While these types of software issues should be resolved by IT before the patches and upgrades are widely deployed, there will always be issues that cause a failure in a small subset of clients—and having the right backup solution in place can mitigate many of the headaches these problems would normally cause.

Support for the computing platforms driven by business needs

Too many environments today contain both PC and Mac clients to commit to a backup solution that requires completely different applications for each client platform. Yet that is often what happens when Mac clients are introduced into the predominantly PC client business computing world. Introducing a completely different application might have no effect on the end user and his computing platform, but it forces an additional support burden on IT because it becomes necessary to spend the time and effort learning a unique application for that specific

If you're relying on network file share or consumer-grade backup for your enterprise desktops and laptops, think again. You're treading on dangerous ground.

Stop treading lightly on dangerous footing—and start enjoying the peace of mind that can only come with using the undisputed, number-one service for automatic desktop and laptop data protection and reliable data recovery. Learn more about the enterprise-class Connected® Backup solution – cloud-based, licensed and hybrid backup and recovery for desktops and laptops. Email digital-info@ironmountain.com or visit www.ironmountain.com/cloudhelp to schedule an assessment of your information management needs.

IRON MOUNTAIN®

environment. Additional overhead for IT is rarely the hallmark of proper product selection.

Roadblocks to Implementation

If the problems related to backup are so easily identified, why hasn't IT everywhere already solved them and moved on to other issues? The simple answer is that backup is a moving target; technologies change and it's the IT department's responsibility to stay ahead, or at least abreast, of the changes to client computing and backup technologies that will make their job easier. But there are a number of common roadblocks to the successful adoption and deployment of corporate-wide backup solutions.

Increase Usability

Usability, from the perspective of the IT department, translates into "how much work will it be to use this solution?" Simplified deployment is the first step to placing a backup solution on every desktop. The question then becomes one of integration with existing deployment tools, how much end-user interaction is required, and how many man hours of IT time is required to successfully deploy and test the backup solution across the enterprise.

Giving end-users the ability to restore their own files or, more accurately, giving IT the ability to delegate file recovery capabilities to business departments or specific users, will reduce the workload on IT and give users a greater sense of involvement and awareness of the need for backup. Simplifying the entire procedure to restore that accidentally deleted file also improves the business workflow, by reducing or eliminating the lost time while important data is restored to the knowledge worker.

Centralized Management

Managing and monitoring the backup process needs to be a centralized activity in the corporate enterprise. Automated report generation and similar tools can keep appropriate IT managers apprised of the state of their backup infrastructure, but the ability to proactively manage that infrastructure is just as important.

Regardless of the location of end-user client hardware in the enterprise, be it a fixed desktop, remote location, or mobile user, IT needs to be able to determine that the client is properly backed up and that critical data is protected. Centralized management tools with the ability to delegate responsibilities is the most realistic answer in today's distributed computing environments.

Time and Efficiency

Time is the most critical component of the backup solution. There is only so much time available in the day to perform backups. Technology has addressed most of the issues that used to be covered by the concept of the backup window. Once a full backup of a client's data has been made, it's no longer necessary to do more than pass the changed bit information to the backup solution. This has the benefit of reducing the amount of time necessary to complete a backup and greatly reducing the amount of data that needs to be passed across the network.

There was a point in history when a business network would take a major hit, because during the allotted time for the network-based backup to run bandwidth availability slowed other network activities to a crawl. Network bandwidth and the performance of backup devices continues to increase and address this weak point in the backup process, but limiting the amount of data, using techniques such as only copying changes at the bit level to the backup, do the most in improving the backup experience.

But nothing is better at improving backup efficiency and reducing the amount of time and poorly utilized storage dedicated to backup, than simply not keeping unnecessary data, especially in large corporate networks. This is where the concept of data deduplication really shines.

Data deduplication allows for more efficient backup storage utilization by reducing the number of copies of a file that is being stored. In any large corporate enterprise, there are always files that manage to duplicate themselves across the network: presentations sent to multiple users, corporate memos, spreadsheets, images, media files, and documents that are passed around in workgroups or across the entire organization and saved by individual users. With deduplication, rather than a thousand copies of that corporate memo stored as backup data, only a single copy per physical backup location may be retained, with pointers in place to restore that copy should any specific user need to recover it. The importance of this technology cannot be understated; it's key to any large-scale efficiency in the backup process.

Security

Concerns over data security are ongoing for IT. The need to know where data is being stored, the security of the transport of that data, and the ability to control data access are daily worries in most IT departments. Security concerns become even more pronounced when data is in more locations than just the primary business site. With the prevalence of remote offices and mobile users utilizing the Internet to connect back to the corporate datacenter, technologies such as data encryption become more than just good to have; they become essential to ensure the security of proprietary corporate data.

Utilizing the Cloud, Public and Private

Making use of the right cloud–based storage and backup service addresses the concerns of today's IT client backup need by solving the most commonly associated problems with client backup. With the installation of a client agent, backup becomes completely automatic. Cloud services generally offer a web-based user portal that allows individual users to recover their own files, without the need for IT intervention. And to assuage security concerns, top-tier solutions encrypt the

data they are backing up at the point it hits the wire; never unencrypting the data while it's under the direct control or storage of the cloud provider.

The cloud doesn't care where the client resides, making it suitable for backing up desktops at any geographic location with a public Internet connection that the company chooses to use. Mobile users are able to back up and restore from their home or hotel room, without the need to make a secure connection back to the corporate network for the sole purpose of backup. An automated client detects the available connection and performs a secure backup whenever the mobile user is connected to the Internet, removing the human element from the backup equation.

Top-tier vendors currently support both the public and private cloud model, with the private cloud being a product installation that resides entirely within the corporate network and not connected to the outside world. But for a company with remote offices and a large mobile workforce, a hybrid scenario—where both the public and private cloud are utilized—offers the highest level of backup availability. A secure and efficient vendor infrastructure is required to support this hybrid solution, but it offers the benefit of a single "throat-to-choke" approach to all of your client backup needs, as well as a common management and user model across your enterprise.

The cloud storage licensing model offers significant flexibility in implementation, easily manageable cost structures, and a high level of availability. The location independence of the client, and the easy access to backup facility via the Internet, means that there is little IT concern that data isn't being protected.

Selecting a Provider

The choice of a provider for your client backup services should be a straightforward one. The provider must meet your immediate business needs and be able to grow with your business, both in the expansion of the number of clients and the adoption of the most effective technologies to address your backup needs. If your business requires it, the provider must also be able to address any regulatory compliance issues that the business may face; not as an external add-on, but as an integrated capability within the product or service itself.

Client backup is a critical component of the IT service delivery process. It needs to be delivered in an effective fashion that minimizes the impact on users and the IT department that supports it. The right enterprise backup solution will support your business without requiring modification of the business workflow process to meet the needs of the backup solution. And, ideally, you can deploy the same solution throughout the enterprise, regardless of client operating system or physical location of the client computer.

Backup as a component of the eDiscovery and compliance process

PCs, laptops, and desktop data have become a tremendous asset and liability for organizations. As the standard hard disk has gotten significantly larger, the amount of critical corporate data exposed on client systems has increased exponentially. Nowadays, key employees are mobile and most of their laptops and attached devices are unprotected. Using network file share as the primary method of backup is expensive and clumsy.

On top of that, organizations face legal costs and risks because they no longer know where critical data is stored; it's difficult to discover data in the wild and apply legal holds. Too often IT faces fire drills due to exposure, and the pressure to provide timely disclosure is often a challenge. Additionally, organizations face issues around data leakage and compliance, especially when it comes to customer information.

This means IT needs to protect endpoint assets, and quickly find relevant PC data for discovery, regulatory compliance and audits. They must simplify collection of laptop and desktop data for legal holds. It is important that they be able to meet regulatory requirements for all of the systems for which they are responsible. This means being able to assure complete and regular backup, having good reporting tools on the state of the overall backup of the company's client computers, and being able to utilize the information provided by the backup tool for internal reviews and to meet compliance standards.

David Chernicoff is a technology consultant with a focus on the mid-market space, *Windows IT Pro* Senior Contributing Editor, founding Technical Director for PC Week Labs (now eWeek), former Lab Director for *Windows NT/Windows 2000 Magazine* (now *Windows IT Pro*), and formerly Chief Technology Officer for a network management tools ISV. David has been writing computer-related feature and product reviews for more than 20 years and is co-author of a number of operating system books, ranging from the *Windows NT Workstation: Professional Reference* (New Riders Publishing), to the *Microsoft Windows XP Power Toolkit* (Microsoft Press), as well as over a dozen eBooks on topics ranging from network switching topologies to production FAX technology.

Having trouble deciding whether to migrate to the cloud?

Iron Mountain can help. We offer a rich set of on-premises, hybrid and cloud information management solutions for archiving, eDiscovery, data protection and compliance. We can help you select the right solutions to meet your specific needs so you can migrate to the cloud when and if you're ready.

Iron Mountain has a proven track record of intelligently managing enterprise information and pioneered cloud information management over 15 years ago. Our underground data centers provide an unprecedented level of security, reliability and protection. Trusted by over 90% of the Fortune 1000, with 140,000 clients in more than 40 countries, we go to great lengths to protect our customers' information, their reputation and ours.

Email digital-info@ironmountain.com or visit www.ironmountain.com/cloudhelp to schedule an assessment of your information management needs.

IRON MOUNTAIN®

ScriptLogic PacketTrap IT

There are two distinct types of network monitoring: the nitty gritty of low-level packet capture and analysis, along with the higher-level monitoring made possible by the likes of SNMP and Cisco's NetFlow. ScriptLogic's goal is to combine these two levels into a comprehensive—though lightweight—package. **PacketTrap IT** pursues this goal with the rather catchy mantra of *Monitor, Alert, Report, Remediate*.

The product's 15-minute initial deployment claim is well founded. Installation proceeds with little fuss, having modest requirements of Windows Server 2003 or later, 2GB of RAM, and a 2GHz CPU for as many as 250 devices.

After an initial device scan, a set of filters and displays assist in the identification of detected devices and, more important, which probes those devices are responding to (e.g., Ping, SNMP v1-3, WMI). Most devices are automatically slotted into groups according to a *policy*, which is a set of predefined probes and alerts that are determined to suit each type of device.

You can deploy agents to enable remote monitoring on an individual basis or in a hierarchical fashion. The latter option builds in some handy resiliency against Internet issues. The local agent will continue to collect stats on devices assigned to it and will reconnect with the main server when connectivity is restored. All this functionality helps make device addition particularly easy and straightforward.

After the initial device scan, the product imports devices into the Devices tab, which provides basic information tables and single-level drill-downs for details. This tab also offers a network-mapping feature that is quick and interesting to view but isn't terribly useful. In future versions of PacketTrap IT, this feature could definitely use some beefing up.

My primary complaint about the Devices tab is its litany of listed tools. These tools are free downloads—not included in the package—that will enhance your toolkit. Unfortunately, most users might skip over the added functionality because PacketTrap triggers an error message when you click a listed tool, referencing the system's lack of licensed software—which, in turn, further confuses the issue. Although I appreciate the desire to reduce bloat, the inclusion of these tools would greatly flesh out the PacketTrap experience. They're already free and readily available, so why not include them in the package by default and save customers the added hassle of another download and licensing gauntlet?

Reporting is a simplified task in PacketTrap. Reports are organized into the top offenders in each category, which range from VoIP call paths to NetFlow/JFlow/SFlows to WMI attributes. In its agent software, PacketTrap includes its own NetFlow generator, which it tweaks slightly to create what it calls *ptFlows*. These ptFlows are fed back to the central server over the local network or via the agent's ability to maintain connections with the PacketTrap server over the public Internet (as opposed to a VPN).

In keeping with this overall notion of simplicity, the dashboard offers a real-time view of similar data in many areas including ,surprisingly, virtual machine (VM) status monitoring at the hypervisor level. A solid, typical bundle of display widgets, meters, charts, and graphs are available as well.

When you access the Administration tab, you'll notice two unique features. First, PacketTrap can perform an automated statistical analysis to establish baselines for your network. Doing so helps PacketTrap avoid notifying you unnecessarily about conditions you already know exist. Second, there's a large management information base (MIB) library available for download to augment its SNMP engine, but this feature can't import custom MIBs—an ever-present concern of many mainstream programs.

PacketTrap boasts a few features that, strangely, didn't make an appearance in my testing. After several hours of experimenting with the product, making every attempt to at least access every feature, I found, for example, no sign of the SNMP trap or router configuration backup interfaces. Combined with the aforementioned "missing" tools in the Devices tab, PacketTrap gives the perception that its design lacks completeness, particularly given its determination to be an all-in-one solution.

PacketTrap IT places itself in the same league as packages from SolarWinds and Ipswitch. Compared with these products, it performs quite well, offering an honest set of primary features with some interesting twists. PacketTrap's ability to leverage the potential of agent-based monitoring is also impressive.

The product's minor flaws can be ironed out; they merely irk users like me who look for context menus to provide logical leaps away from the current display. Nonetheless, PacketTrap will work for most organizations looking for a good core performer with forward-thinking innovations under the hood.

InstantDoc ID 125944

ScriptLogic PacketTrap IT

PROS: Simple, stable, lightweight monitoring that extends agent-based monitoring beyond its traditional roles; compares admirably with others in its class

CONS: Some features seem extraneous or under-developed; the product's simplicity sometimes makes features seem isolated because it lacks right-click interweave

RATING: ◆◆◆◆◇

PRICE: $1,995 for 50 devices; licensing is per device

RECOMMENDATION: If your monitoring and alerting needs are modest, I recommend PacketTrap IT, which takes on the challenge of implementing or facilitating initial corrective steps. However, the product needs a bit more focus and cohesion.

CONTACT: ScriptLogic • 561-886-2400 • www.scriptlogic.com

Brandon Carse | bcbigb@gmail.com

REVIEW

Exclaimer Auto Responder

Automated email replies are the heart of **Exclaimer Auto Responder**, but it also does a lot more. It can redirect messages, modify senders, add recipients to mail, and implement ethical walls between groups. Microsoft Exchange Server has some of these features, especially Exchange 2010, but Auto Responder provides them with granular control and easy manipulation of automatically generated messages.

Auto Responder's installer is simple and takes less than ten minutes. It comes as a single .exe file that's compatible with Exchange 2007 SP1 Rollup 5 and later and Exchange 2010. It's 64-bit–only, in keeping with the needs of Exchange, and runs on Windows Server 2003 or later Windows Server OSs. Auto Responder is deeply tied into Exchange, operating as a custom transport agent, so it needs to be on an Exchange server with the Hub Transport role.

The final part of setup is configuring your first policy. The New Auto Responder Policy wizard is well laid out and offers relevant information for all the available policies. I have one small complaint about the wizard: Once you've completed the wizard, you must click save in the main console or your new policy won't take effect. I'd rather have things work immediately after finishing the wizard.

With setup being so simple, you can focus on understanding the product. You can apply a huge amount of configuration to policies—Figure 1 shows some of the possibilities. The settings will be very familiar if you've ever set up an Exchange Transport Rule. There are far too many possibilities to list them all here, but some key examples include close control of the subject of the returned message, the ability to covertly add recipients, granular control over who policies apply to (such as being inside or outside the organization or a member of a group), presenting a different external email address from your main company address (which can be great in a merger scenario), preventing groups from sending mail to one another, providing a specific response only during certain time periods (like an out of office message), and allowing a catchall policy for any particular domain.

Auto Responder doesn't simply reroute or add a recipient to messages that require an automatic response. Some of my favorite options include adding attachments to an auto response and adding the sent message to the reply just like Outlook does, to make the response seem more human. The key benefit of Auto Responder is the control it gives you over the response message. A template library in the product has a few examples of different auto response messages, but it would have been nice to have a whole bunch of different examples to build on.

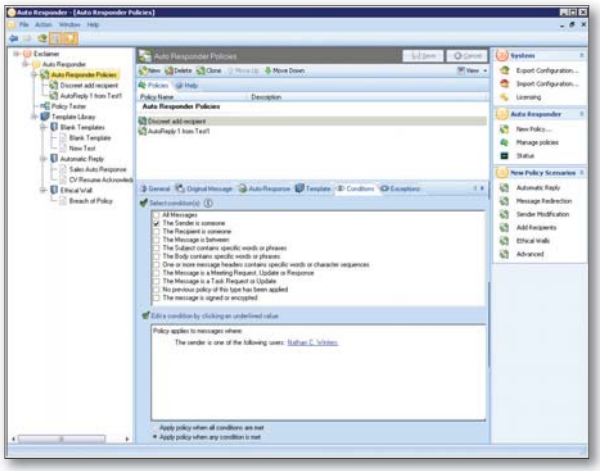

Figure 1: Auto Responder Policies

If you have multiple Exchange Hub Transport servers and want to maintain a single set of policies, templates, and settings across all of them, you have several options. The simplest is to back up your configuration by exporting it to a file, which you can then import on another server. You can also do it automatically by installing Auto Responder on each Hub Transport server and then setting up a central file share. There's a security question to note with this method: The Everyone group must be given read permissions to this share because the service the software runs under is Local System. These permissions could allow anyone to discover the share and view potentially sensitive policies within. Exclaimer suggests using a hidden share, but security by obscurity isn't a great idea in my opinion.

Other than the possible security problem, the process works well. When administrators make any changes, they're prompted to apply them on the remote servers. The changes are saved to the shared folder, which the software monitors for changes. Changes are applied within about 10-20 seconds.

Exclaimer Auto Responder is a well thought out, well documented, clear, and effective product that gives administrators the ability to automate a variety of actions to be taken based around almost infinite scenarios. Notification and response templates ensure that a company brand can be maintained throughout your communications.

InstantDoc ID 126052

Auto Responder

PROS: Quality documentation; very granular policy settings; easy to configure

CONS: Questionable security with remote deployment feature; irritating need to save settings after completion of policy wizard; could use more response templates out of the box

RATING: ◆◆◆◆◇

PRICE: $495 per server and $99 per year for support

RECOMMENDATION: Exclaimer Auto Responder is a well thought out, comprehensive, and effective product that gives administrators the ability to automate a variety of actions in many scenarios. The product helps maintain control and branding, features likely to be most interesting to large organizations.

CONTACT: Exclaimer • 888-450-9631 • www.exclaimer.com

Nathan Winters | nathan@clarinathan.co.uk

3 Network-Monitoring Systems

Automate your network monitoring with one of these comprehensive products

by Nate McAlmond

At the midsized company I work for, we've been using an older network-monitoring system for about seven years. It gives our administrators basic up or down information about servers and services, and it sends SMS (text message) and phone alerts to our cell phones in the event of problems. I decided it was time to upgrade or at least add a more current tool that could provide improved performance and detailed status information about our Exchange Server, SQL Server, and Terminal Server systems on the network at an affordable price. I came up with three viable candidates: Ipswitch's WhatsUp Gold Premium, ManageEngine's OpManager Professional, and SolarWinds' ipMonitor. All three of these network monitors cost less than $3,000 (for 100 devices) and are available for trial periods so that you can test drive them for yourself. Let's see how they stack up.

The Discovery Process

To prepare for testing, I first needed to enable SNMP on all devices, including our Windows servers. I set up read-only SNMP access by configuring the SNMP settings for each device on the network that I wanted to monitor.

In Windows Server 2003/2000, you install the service from Add/Remove Programs, Windows Components; in Windows Server 2008, you add the SNMP feature from Server Manager. Then, you need to go into the Control Panel Services applet and configure the SNMP service—a quick and easy process. Managed network devices such as firewalls, switches, routers, and printers will also have SNMP management capability and are usually quite easy to configure. For more information about SNMP, see the Microsoft appendix "Simple Network Management Protocol" (technet.microsoft.com/en-us/library/bb726987.aspx).

Next, I installed all three monitoring systems on one of two Windows XP SP3 machines at my desk. Once installed, each system consisted of a database and web server. In all three monitoring systems, the system is meant to be used from the web server interface by multiple individuals, and you can configure credentials with full or varying levels of limited access. Each user at either system has the ability to add, remove, and reposition widgets on his or her dashboard. Widgets give insight to a particular aspect, such as processor or memory utilization, across many devices on the network.

Before using each system to scan the network (called the *discovery process*), I added credentials that the system would use to gain access to each device that it discovered on the network. As you can see in Table 1's comparison chart, Ipswitch's WhatsUp Gold Premium has credential options for SNMP, WMI, Telnet, SSH, ADO, and VMware. ManageEngine's OpManager Professional has options for SNMP, WMI, SSH, Telnet, and URL. SolarWinds' ipMonitor has options for SNMP, WMI, and URL.

After configuring SNMP at the network devices and credentials (Windows and SNMP) for each of the three network-monitoring systems, I ran the discovery process on a range of IP addresses in my local subnet. All three systems discovered about 70 devices. Using the default scanning options, the three monitoring systems did a good job identifying each device type, as well as providing a lot of insight into device status. All three included several monitors for what you would expect in a network-monitoring system, such as processor, memory, disk usage/utilization, packet loss/latency, Exchange/Lotus, Active Directory, and every Windows service on the system. All three products had the ability to add monitors to each device one at a time or to large groups of devices all at once.

Both OpManager and WhatsUp Gold include interfaces for identifying and collecting VMware events on VMware host and guest servers. Also, both OpManager and WhatsUp Gold contain a switch port mapper, which shows you what device is connected to each port on your managed switches. With this information, you'll know which port on each switch contains the link to critical business applications without having to physically trace wires in your server rooms. Then, you can configure alerts on the switch devices for these individual ports. With OpManager, it's easy to get the switch port mapper results by selecting the switch and running the Switch Port Mapper tool; the product returns results within a few seconds. In the case of WhatsUp Gold, the tool is called MAC Address and needs to be run with the *Get connectivity* option turned on. WhatsUp Gold takes much longer to get the results, as it appears to be scanning and calculating connectivity information for the entire subnet.

Ipswitch WhatsUp Gold Premium

PROS: Provides the most accurate information of all three products; lets you create custom monitors; offers comprehensive VMware monitoring tools; integrates with AD

CONS: Not as many built-in monitors as the other products; more expensive than the other two (until you get up to 500 nodes)

RATING: ◆◆◆◆◆

PRICE: $7,495 for 500 devices; $2,695 for 100 devices; $2,195 for 25 devices

RECOMMENDATION: I recommend WhatsUp Gold for IT shops that have a large VMware installation or a desire to create their own monitors

CONTACT: Ipswitch • 800-793-4825 or 781-676-5700 • www.ipswitch.com

Ipswitch WhatsUp Gold Premium

With both ipMonitor and OpManager, I would occasionally notice strange readings that made me wonder what was going on. With ipMonitor, when processor utilization reached a low percentage, it would sometimes show as a negative value on the dashboard widgets. In another instance of low processor utilization, ipMonitor sent me an alert that a processor was at 11,490 percent utilization! OpManager would sometimes track and send me accurate information about domain controller (DC) disk utilization but not show it in the top 10 disk utilization widgets on the dashboard—but right above that widget was another widget showing me that one of our DCs should be in the top three. No such situations occurred when I used WhatsUp Gold. In fact, WhatsUp Gold tracks processor cores in its processor utilization widgets, and when I compared the WhatsUp Gold processor widget reading with the Windows Performance Monitor results, it always appeared accurate on all individual cores. Likewise, hard disk utilization was accurately reported on all hard disk dashboard widgets.

WhatsUp Gold includes a monitoring library, which lets you create new monitors based on other monitors. Larger organizations might find this feature useful for ensuring a uniform set of monitoring by device type across a large technology environment; this seems the most efficient way to adjust monitors on groups of devices.

WhatsUp Gold doesn't contain vendor-specific monitors (except for an APC UPS monitor)—as OpManager does for Dell, HP, and IBM—but rather contains an option for Active Script monitors, letting you use VBScript or JScript to write your own monitoring routines. There's also an online resource center for Active Script, where the WhatsUp Gold user community can download and share code.

One improvement I'd like to see WhatsUp Gold incorporate would involve the UI—which Figure 1 shows—mainly because it's so linear. For example, after drilling down into the Active Monitor Library, it takes about five Cancel or Close clicks to return to the dashboard. Also, unless scripted, WhatsUp Gold doesn't have a monitor for checking the status of a website, which might be necessary—especially if it's hosted on somebody else's server that you don't have another way of accessing.

For situations in which a device stays down, you can configure WhatsUp Gold to send alerts at 2, 5, or 20 minutes. So, it's possible to escalate the alert in situations when the expected individual doesn't respond within that time frame.

Of the three tools, WhatsUp Gold is the only one with the option to integrate with an LDAP environment, which could be particularly valuable to large deployments.

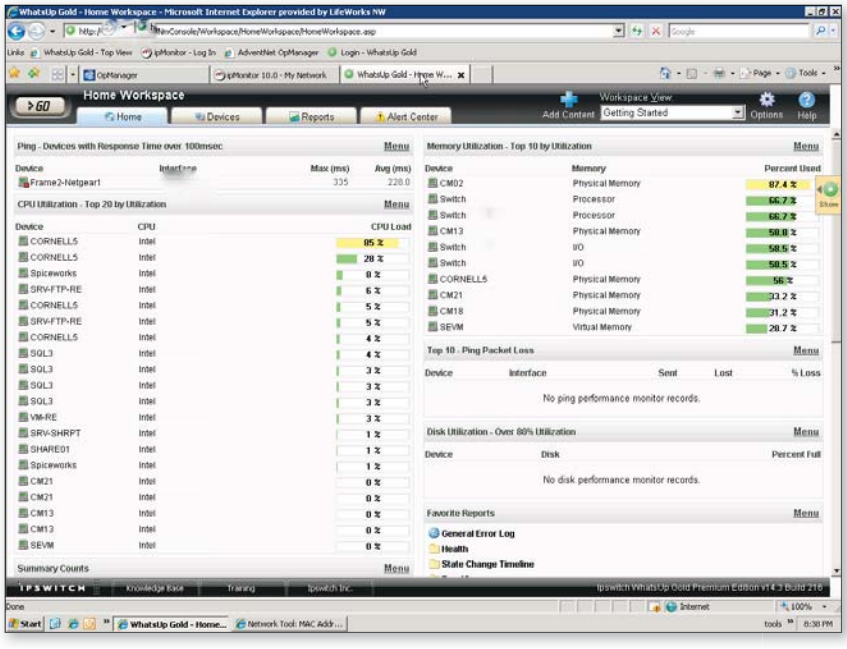

Figure 1: The WhatsUp Gold Premium UI

ManageEngine OpManager Professional

PROS: Best overall UI; more built-in monitors than the other two products; lowest pricing for 50 or fewer nodes

CONS: Didn't monitor all devices accurately in my tests, so it might require some troubleshooting to get it completely functional

RATING: ◆◆◆◆◆

PRICE: $1,995 for 100 devices; $995 for 50 devices; $595 for 25 devices

RECOMMENDATION: IT shops looking for the maximum amount of built-in functionality (with the exception of AD integration) will appreciate OpManager Professional. In the 26- to 50-device range, it's about half the price of the other two products.

CONTACT: ManageEngine • 888-720-9500 or 925-924-9500 • www.manageengine.com

ManageEngine OpManager

Upon installing OpManager, I noticed that it was simple to navigate and configure a very wide range of features. OpManager has the option (along with SMS and email) of sending a Direct Message to a Twitter account—a nice alternative to email. Using the Twitter account this way gives me

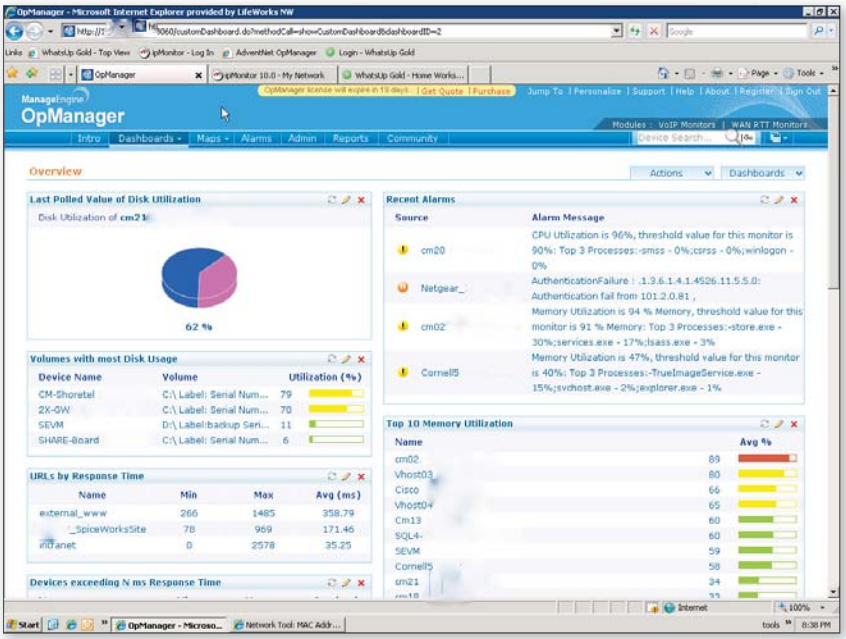

Figure 2: The OpManager Professional UI

an overview of what's happening on the network, but because my phone doesn't ring when a tweet comes in, I still want the text-message alert for anything critical. I can see myself setting low thresholds for a Twitter alert on all server devices, just to get a good running log of what's going on, but not necessarily using it for warnings about critical situations.

In addition to its expected monitors, OpManager offers vendor-specific SNMP performance monitoring for Dell PowerEdge, HP ProLiant, and IBM Blade Center devices. OpManager can also integrate with the Google Maps API so that you can add your devices to a Google map. However, you'll need to pay for a Google Maps API Premier account (unless you plan to make your network map publicly available), as this is a licensing requirement for the free version of Google Maps API.

For situations in which an administrator receives an alert but doesn't respond within a certain amount of time, you can configure OpManager to send another alert to a different administrator. For example, for an individual who might be out sick or distracted but would normally respond to a critical event on a particular group of servers, you could have an escalation alert notify another administrator if the original alert isn't acknowledged or cleared within a specified number of hours/minutes.

OpManager was also the only one of the three products to include a section specifically for monitoring the quality of VoIP over the WAN. The VoIP monitoring tools in OpManager require devices at both the source and destination networks to support Cisco IP SLAs. In addition, OpManager—whose UI you can see in Figure 2—includes more monitors and dashboard widgets than either of the other two products.

SolarWinds ipMonitor

PROS: Unlimited devices at a very low price; easy to use

CONS: No multiple-administrator communication mechanism

RATING: ◆◆◆◆◇

PRICE: $1,995 for unlimited devices (25 monitors free)

RECOMMENDATION: If you're on a tight budget and you need to monitor a lot of devices, or if your monitoring needs aren't terribly complex and you're OK using an outside-the-system way of communicating between administrators, SolarWinds is for you.

CONTACT: SolarWinds • 866-530-8100 or 512-682-9301 • www.solarwinds.com

SolarWinds ipMonitor

When I first installed ipMonitor, I found the UI—which Figure 3 shows—confusing. For some reason, it took me forever to find the setting for how often the product would check individual system monitors (which turned out to be every 300 seconds). However, after using ipMonitor for several weeks, I found it to be extremely easy to use, offering enough monitoring options to keep a good eye on the network. With ipMonitor, you can adjust the default scan so that any service or performance option will always

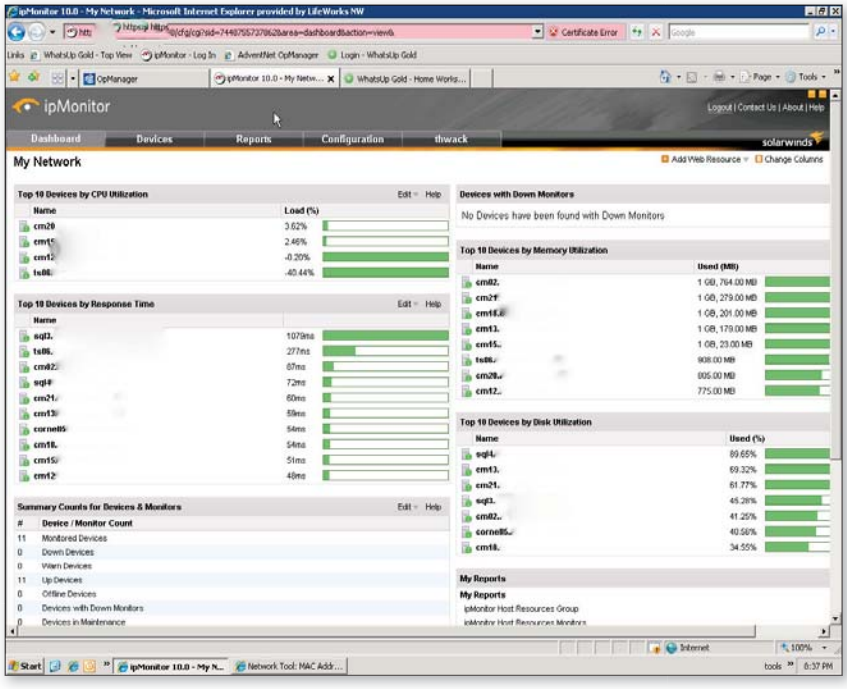

Figure 3: The ipMonitor UI

3 NETWORK-MONITORING SYSTEMS

Table 1: Comparison Chart

	Ipswitch WhatsUp Gold Premium	ManageEngine OpManager Professional	SolarWinds ipMonitor
Credential Options	SNMP, WMI, Telnet, SSH, ADO, and VMware	SNMP, WMI, SSH, Telnet, and URL	SNMP, WMI, and URL
Database and Web Server	MS-SQL; proprietary web server or IIS	MySQL; Apache Tomcat	Proprietary for both
Alerting Options	Email, modem, SMS, restart service, log event, Windows pop-up, SNMP Set, VMware command, SSH, Syslog, Active Script, and more	Email, GSM modem, SMS, Twitter Direct Message, run a program, and run a system command	Email, Net Send, modem, SMS, log event, external process, SNMP trap, restart server or service, and more
Windows Event Log Scan	Run a report	Includes monitoring of event logs and can alert specific event types	Includes monitoring of event logs and can alert specific event types
Maintenance Intervals	Yes	Yes	Yes
SNMP Active/ Passive	Yes	Yes	Yes
User Rights	Most granular control	Least granular control	Granular control
Escalation Alerts	Yes, maximum window of 20 minutes	Yes, most flexible	No
URL Monitoring	No	Yes	Yes
SSL	Yes	Capable, but you must configure this within the Apache Tomcat web server	Yes
Virtual Event Collecting	VMware	VMware	No
Switch Port Mapping	Yes	Yes	No
LDAP Integration	Yes	No	No
WAN Latency	Ping	Cisco IP SLA or Ping	Ping
VoIP Monitoring	No	Cisco IP SLA	No
Multiple-Administrator Communication Interface	Yes	Yes	No
Map View	Layout only	Layout with image; Google Maps API	Layout with image

be selected in future scans. In addition to the expected (and previously mentioned) monitors, ipMonitor offers a Windows event log monitor that you can configure to check for events that you want to be alerted about.

On the downside, ipMonitor doesn't contain any type of alert tracking/assignment system. If you have only one network administrator, this should be no big deal, but larger IT departments will probably find that the inability to acknowledge/assign/clear alerts is a problem. Unless the administrators remember to communicate outside the system, multiple administrators could receive the same alert and start working on the same problem. Perhaps devising an agreed-upon way of responding—for example, if you get the alert, it's your responsibility to take care of it, or dividing up the network devices so there's no question about who should respond—would do enough to prevent this situation.

Decision Time

I've just made my decision about which of the three products will best suit my environment. I'm going with ManageEngine OpManager, licensed for 50 devices, for a few reasons.

First and foremost, I want to be able to monitor as many aspects of our environment as possible because I realize it's the best way to avoid unexpected outages. OpManager definitely offers the most breadth of all three products in this area. The second reason is budget. I can continue to use our existing up-or-down monitoring tools for our workstations and printers, so I can avoid paying for the additional licensing. Third, I really like the way the people at ManageEngine have designed OpManager to take advantage of new technology, and I feel it's worth paying the annual service and support fee to get the updates as this product develops.

InstantDoc ID 125832

Nate McAlmond
(mcalmond@gmail.com) is an MCSE, a Security+, and Network+. He's the director of IT for a growing social services agency, and he specializes in thin-client infrastructure and electronic health record implementation.

Exchange Server Backup and Recovery Software

Pick the right solution and you can forget all about backup troubles

by B. K. Winstead

Backup: the most important thing in IT that everybody loves to hate. And it's not just the backup process itself; after you've made that backup, you need to be sure you can recover data from it. When you're talking about your Microsoft Exchange Server, recovery might mean anything from one accidentally deleted email message to a mailbox database to a whole server.

Fortunately, there's a plethora of third-party products that can help you automate this process. Furthermore, Exchange Server 2010 has changed the landscape and requirements for backup software through the introduction of database availability groups (DAGs), so you have something else to consider if you've made the move—or are planning a move—to the latest Microsoft mail server. The accompanying table gives you a snapshot of some of the available products and should help you start your search. Let's review some of the key product differentiators you might want to keep in mind.

Big Differences

The first question you might ask is whether you need a dedicated Exchange backup product or something that can do global backups of your IT infrastructure but is also Exchange-aware. You'll find plenty of both options available. If you go with the larger product, the Exchange piece might need to be purchased as a separate add-on or module. Just remember, the added backup functionality typically brings with it added cost and complexity.

Perhaps the next consideration is what method of backup best suits your organization. Methods run from from brick-level full server backups to byte-level replicated changes. Keep in mind that how backups are taken can potentially affect the granularity of restore. Make sure your current infrastructure includes the necessary hardware and software prerequisites, or you'll need to factor in additional upgrades as you calculate your potential costs.

Related to how things are backed up is how easily data can be restored. Obviously, any backup is worthless if you can't pull from it what you need when you need it. If you're hoping to keep in the good graces of the C-level exec who constantly deletes that most vital email message, make sure the product you choose will be able to quickly and easily locate and restore single mail items. Some products also work at or near continuous data protection (CDP), which means minimal loss of data when disaster strikes.

Finally, keep in mind the target storage medium. Many products give you options—disk, tape, even cloud storage these days. As always with backups, it does you little good to keep your backup copy in close physical proximity to the original data—you don't want to live with the possibility that both versions could be damaged or destroyed at the same time.

Exchange 2010 DAGs

The DAG is probably the biggest architectural change that came with the Exchange 2010 release. It lets you set up copies of the primary mailbox database such that the system automatically fails over to a copy if the currently active database encounters a problem. You can have as many as 16 copies of each mailbox database.

DAGs require you to run Exchange on the Enterprise edition of Windows Server 2008 or Server 2008 R2. You also have to consider the cost of additional server hardware and licensing; depending on the size of your organization and the number of DAG copies you choose to create, the cost could end up being quite high. However, a well-designed DAG architecture with at least three copies of every database could eliminate the need for a third-party backup product. Microsoft claims to have been running without traditional backup since it implemented Exchange 2010 internally during the product's development.

You might want to implement DAGs and use a third-party backup product as well. The only requirement to keep in mind for such arrangements is that the backup solution must use Microsoft Volume Shadow Copy Services (VSS) for backup; streaming backups won't work with DAGs.

The Total Package

The accompanying table isn't comprehensive, but it should give you a start. There are many other products out there. Additionally, this guide doesn't address cloud-based solutions, and there are many products in that space. So, do your research, and be sure to pick the product that offers the total package that works best for your organization. And don't forget to test those backups once in a while—no one wants to find out their backup files are no good when they really need them.

InstantDoc ID 126058

B. K. WINSTEAD
(bwinstead@windowsitpro.com) is an associate editor for *Windows IT Pro* and *SQL Server Magazine*, specializing in messaging, mobility, and unified communications.

Company	Product	Price	Exchange Versions Supported	Dedicated Exchange Product or Add-On to Another Product?	Installs on What OSs?	Other Hardware/ Software Requirements?
Acronis 781-782-9000 877-669-9749 www.acronis.com	Acronis Recovery for Microsoft Exchange	$1,219	2007/2003/2000	Dedicated	Windows Server 2008/2003/2000, Windows Vista/XP	None
AppAssure Software 703-547-8686 www.appassure.com	Replay	$899 per Windows server	2010/2007/2003	Add-on	Server 2008 R2 (x86 and x64)/2008/2003, Windows 7/Vista/ XP	None
Arkeia Software 760-431-1319 www.arkeia.com	Arkeia Network Backup	Starting at $2,500	2010/2007/2003/2000	Add-on	Backup Server on Linux, VMware VM, UNIX	None
Atempo 650-494-2600 www.atempo.com	Time Navigator	Contact vendor	2010/2007/2003/2000	Add-on	All the major OSs	None
Axcient 800-715-2339 www.axcient.com	Axcient Complete Data Protection Solution	Varies based on sizing; contact vendor	2010/2007/2003	Dedicated	Windows, Mac, Linux	None
Lucid8 425-456-8473 www.lucid8.com	Exchange Protection Manager	$599 per Exchange server; unlimited mailboxes	2010/2007/2003	Dedicated	Server 2008 R2/ 2008/2003	Client OS: Server 2008 R2/ 2008/2003, Windows 7/ Vista/XP
NetApp 877-263-8277 www.netapp.com/us/	SnapManager for Exchange	Contact vendor	2010/2007/2003	Add-on	Server 2008 R2/ 2008/2003/ 2000	NetApp Snap Drive 6.1 or later, NetApp SnapRestore software, FCP or iSCSI protocol, Microsoft Management Console, Windows PowerShell 1.0
Novosoft 383-330-3469 www.novosoft.net	Handy Backup	See www .handybackup .net/volume-pricing.shtml	2010/2007/2003	Add-on	Windows 7/ Vista/XP	None
UltraBac Software 425-644-6000 866-554-8562 www.ultrabac.com	UltraBac Exchange Agent	$930, including one year of maintenance	2007/2003/2000/5.5	Add-on	Server 2008 R2/ 2008/2003/2000, and corresponding workstation OSs	None
Vision Solutions 317-598-0185 www.doubletake.com	Double-Take Availability	$3,295 per server	2010/2007/2003	Dedicated	Server 2008/ 2003/2000	None
Vision Solutions 949-253-6500 800-683-4667 www.visionsolutions.com	Double-Take RecoverNow	$1,295 per agent; $1,995 per Exchange agent; $1,995 per repository	2007/2003/2000/5.5	Add-on	Server 2008/2003	None

Editor's Note: Information in this buyer's guide comes from vendor representatives and resources and is meant to jumpstart, not replace, your own research; also, some products might have been left out, either as an oversight or from lack of vendor response.

	Backup Storage Medium Options	Backup Method	Data Dedupli-cation?	Exchange Log Checking?	Supports PowerShell?	Integrates with System Center Operations Manager?	Integrates with Native Exchange Replication?	Can Be Used for Failover?	Recovery by?
	Disk, tape, removable disk, optical/write once, online/cloud	Snapshot	No	Yes	Yes	No	No	No	Server, database, storage group, mailbox, individual message
	Disk, removable disk	Snapshot, replication	Yes	Yes	No	No	No	Yes	Server, database, storage group, mailbox, individual message
	Disk, tape, removable disk, online/cloud	Full, copy, incremental, differential	Yes	Yes	No	No	No	No	Server, database, storage group
	Disk, tape, optical/ write once	VSS snapshot, native mode	Yes	Yes	Yes	No	Yes	Yes	Server, database, storage group, mailbox, individual message
	Disk, online/cloud	File backup, Exchange brick-level, server backup (image backup and snapshot)	Yes	Yes	No	No	Yes	Yes	Server, database, storage group, mailbox, individual message
	Disk	VSS snapshots	No	Yes	No	Yes	Yes	No	Database, storage group, mailbox, individual message
	Tape, online/cloud	Snapshot	Yes	Yes	Yes	No	No	Yes	Server, database, storage group, mailbox, individual message
	Disk, removable disk, optical/write once, online/cloud	Full online backup	No	No	Yes	No	No	Yes	Database, storage group, mailbox, individual message
	Disk, tape, removable disk, optical/write once, online/cloud	Streaming backup	No	Yes	Yes	No	No	No	Server, database, storage group, mailbox
	Disk	Byte-level replication	No	No	Yes	Yes	No	Yes	Server, database, storage group, mailbox
	Disk	Replication with target-side snapshots	Yes	No	No	No	No	No	Server, database, storage group, mailbox, individual message

■ Mobility ■ Disaster Recovery ■ Certification

INSIGHTS FROM THE INDUSTRY

Management for Employee-Owned Mobile Devices

A lot has been written about whether Microsoft's Windows Phone 7 will have any impact on the smartphone market by the time devices start hitting the stores. At one point, Windows Mobile was a leading mobile platform, particularly for businesses, but Microsoft failed to move quickly with the changing times as Apple launched the iPhone and revolutionized the smartphone market.

I understand why a behemoth of a company such as Microsoft is less agile, particularly when they already have (or had) a successful product in the space. It can be very difficult to abandon what has worked—until it proves to be no longer viable. This is why I find what Zenprise is doing these days particularly impressive.

Responding to customer requests several years ago, Zenprise added support for Blackberry devices. And according to Ahmed Datoo, Zenprise's vice president of marketing, "We have expanded our product actually now to support in terms of smartphones, Blackberrys, iPhones, iPads, Android devices, Windows Mobile, and Palm devices—so, typically all the major types of platforms that you'd see in the enterprise."

Zenprise MobileManager has evolved as a complete end-to-end mobile management platform, and now the company is adding features to address the growing number of employee-owned devices connecting to corporate systems. The new additions, available as part of the latest release of MobileManager released this week, are Profiles, Remote Lifeline, and Selective Wipe.

The Profiles features lets IT departments establish different profiles for corporate-owned versus employee-owned devices. So, on corporate assets, you can apply that tighter level of security policies that on employee-owned devices would cause an uproar. It's all still managed through the Zenprise console and tied into Active Directory.

The Remote Lifeline feature lets your IT department support those employee-owned devices that wouldn't normally receive support. Device wipe can be a problem for any business, particularly if you work with sensitive data. The problem gets even more complicated when you allow those employee-owned devices in. Zenprise's new Selective Wipe feature is a pretty good answer. It lets you clear all work-related data from the device without affecting the user's personal data, such as pictures, music, applications, and so forth. Unfortunately, this feature works only on iPhones or iPads because it's based on how Apple tags data that comes from Exchange.

iPhones (and now iPads) continue to make inroads on the corporate landscape. As Jayaram Bhat, CEO for Zenprise, told me, "Today we hear 8 out of 10 prospects ask us about the iPhone. Even those companies who are supposedly standardized on Blackberry want to make sure they have the ability to support iPhones, especially the consumer-owned ones. So, it's a dramatic change—this is a very dynamic market."

The complexity and speed at which the mobility marketplace is changing make it all the more important that IT departments have a handle on what their employees are using—and what they're using them for. Taking a look at a mobile management product such as MobileManager might save a lot of time, money, and headaches in the long run.

Has your IT department developed a strategy for mobile device management, or are you using a third-party product already? Send me a note at bwinstead@windowsitpro.com to let me know.

—B. K. Winstead

How Certification Can Still Be Relevant, a Decade into Your Career

At the end of the '90s, there was a rush toward certification. A large number of people who were breaking into the industry back then used certifications as a crowbar to get in the door at their first job. Many of those people are still in the industry today, working as systems administrators managing servers that support hundreds or thousands of clients.

It is interesting to note that many of these people never took another exam once they settled into their career. After getting their foot in the door, certification no longer seemed relevant. They had too much on their plates to go chasing new certificates when the old ones seemed to have served their purpose.

However, that being said, I'm noticing in my role as an MCT and an author of certification-related textbooks that a lot of people whose last certification was an exam in the NT4 MCSE track are today showing interest in recertifying.

There are good reasons for this. In part it is because they want something a little more up-to-date than an NT4 MCSE on their resume. In part it is because many of them have been in the same role for a while and they want to broaden their knowledge and learn something new. In part it is because they use certification as a set of goalposts to structure their approach to learning a product.

Of course, in theory, they could learn something new by just picking up a book or reading a bunch of whitepapers and articles. So why are these people choosing to return to certification? I have some theories:

- ***Certifications provide a structured framework to learning.*** They give you something to learn and something to measure that knowledge against. It isn't about the designation—these people have been in their career long enough that the actual certification itself won't make much difference to their resume. Instead they are leveraging the certification process as a way of structuring their learning.

It is interesting to note that many of these people never took another exam once they settled into their career.

- ***Certification objectives cover features of a product that you might not be aware of if you just seek knowledge to complete a specific task.*** For example, you might have taught yourself how to use System Center Configuration Manager (SCCM) so that you could use it for operating system and application deployment. You don't bother to learn anything more about the product because you don't have time to dig deeper and you don't have any reason to learn it. If, on the other hand, you were looking toward a certification in SCCM, you'd have to learn about all aspects of the product, not just those that were OS and application deployment related.
- ***Certification gives you a goal.*** Although people might pick up a book on a particular product hoping to learn more about it, most don't get past the first few chapters because there is no structure to their learning. Think about it this way—when you buy a textbook about a product you use, do you read it from cover to cover? Or do you just go to the index, find the page that solves your current problem, and then put the book back on the shelf? Certification gives you a reason to broaden your focus on a product—a reason to learn about those aspects of the product that you don't immediately need. Setting a goal that you can accomplish can make you more motivated to learn than just saying to yourself, "I'll read this book and then pick up another."

Unless you are particularly disciplined as a learner, you'll probably benefit from a structured approach. Having a set of topics that you have to learn and an exam that you can test that learning against gives you a set of goalposts.

—Orin Thomas

Top 5 Considerations for Disaster-Recovery Planning

Below are the top five considerations—according to Nasuni, a cloud-based storage vendor—to keep in mind for effective disaster-recovery planning.

Downtime. A study by Contingency Planning Research estimated the cost of downtime at roughly $18,000/hour for many businesses. How much time it takes to recover data—how long a business can afford to be out of business—is a high priority.

Recovering terabytes of data from tape involves first identifying which set is needed, requesting a delivery from the offsite storage provider, correlating each tape with a logbook, determining which to load first, then restoring data. This process may take days of working round-the-clock, depending on the amount of data. Disk is also vulnerable to corruption and accidental erasures.

Data Integrity. Traditional data restores are often less than 100 percent successful; some files are simply gone for good. When these files pertain to customers, transactions, or anything else not easily reproduced, lost data becomes lost revenue.

As a data repository, tape is notoriously unreliable. With a full day between backups, at least twelve hours of data will be lost. Surveys suggest that up to 20 percent of nightly backups do not successfully copy all data.

Disk mirroring provides data redundancy. In the event of a disaster, however, all data produced since the last backup will still be lost. Synchronous replication does not fully protect against data loss because if software is corrupted or data is deleted from the main server, or in a virus attack, that problem will be copied to backups.

Cost. Tapes themselves are relatively cheap. Disk mirroring essentially requires purchasing a duplicate set of servers for data. As data volumes grow, so do backup costs, often into the hundreds of thousands of dollars, or more.

Simplicity. Traditional strategies effectively drop IT administrators into a maze of hardware and bookkeeping, whether it involves putting a tape in the drive and waiting or piecing data together from corrupted disks. Mirroring disk to a second site can improve reliability but often the task of finding a second data center and establishing a reliable link can be daunting and filled with unexpected complexity.

Security. Just as primary storage must be protected, so must its backups. Most reputable backup facilities provide superb security, except that tapes are sometimes lost or damaged in transit. With disk, the security of data is entirely dependent on internal security mechanisms. If they are secure, your backups will be secure, except for cases of human error or malicious attack.

The recent generation of cloud storage gateways dramatically simplifies disaster-recovery planning. These gateways are designed to bridge your existing operations with the reliability offered by the major cloud storage providers. Gateways are typically packaged as virtual machines (VMs) that are available in all the major hypervisors.

The use of advanced caching algorithms allows data recovery from the cloud to be virtually instantaneous: if a critical server is lost, a quick download of the gateway VM can reestablish the connection to the cloud so that all of the data is available. Leading cloud storage gateways allow customers to meet their recovery-time objectives by prioritizing access to the most critical data first.

Cloud storage providers create multiple copies of data and store them in many disparate servers. Should a server fail, data is already safe in several others. Customers that leverage cloud storage do not need to be concerned with tapes, or drives, or differing systems of inventory. In the event of data loss, recovery is as simple as logging back in to your cloud account.

Cloud storage involves inherent vulnerabilities, but provided all data is encrypted before it is transmitted to the storage service provider, and other common-sense steps are taken, the risks are practically nil.

—Jason Bovberg

Statement of Ownership

Statement of Ownership, Management, and Circulation for *Windows IT Pro Magazine* as required by 39 U.S.C. 3685; *Windows IT Pro Magazine*, publication no. 1552-3136, filed October 1, 2010, to publish twelve monthly issues each year for an annual subscription price of $49.95. The mailing address of the office of publication, the headquarters of general business of Peg Miller, Publisher, and Michele Crockett, Editorial and Custom Strategy Director, is 221 E. 29th St., Loveland, CO 80538. The owner is Penton Media Inc., 249 W. 17th St., 4th Floor, New York, NY 10011-5390. Penton Business Media Holdings, Inc., of 249 W. 17th St., 4th Floor, New York, NY 10011-5390 owns 100% stock in Penton Media, Inc. The average number of copies of each issue published during the twelve months preceding the filing date include: total number of copies (47,874); paid mail subscriptions (30,068); sales through dealers and carriers, street vendors, and counter sales and other non-USPS paid distribution (5,696); paid distribution through other classes of USPS mail (288); total paid circulation (36,052); free or nominal rate distribution by mail (6,766); free or nominal rate distribution outside the mail (2,759); total distribution (45,577); copies not distributed (2,297) for a total of (47,874) copies. The actual number of copies of single issues published nearest to the filing date include: total number of copies (38,391); paid mail subscriptions (19,023); sales through dealers and carriers, street vendors, and counter sales and other non-USPS paid distribution (4,843); paid distribution through other classes of USPS mail (199); total paid circulation (24,065); free or nominal rate distribution by mail (12,259); free or nominal rate distribution outside the mail (592) total distribution (36,916); copies not distributed (1,775) for a total of (38,691) copies.

I certify that the statements made by me above are correct and complete: —Peg Miller, Publisher.

AD INDEX

For detailed information about products in this issue of *Windows IT Pro*, visit the web sites listed below.

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- *asp.netNOW*
- *DevProConnections UPDATE*
- *Exchange & Outlook UPDATE*
- *Security UPDATE*
- *SharepointPro Connections UPDATE*
- *SQL Server Magazine UPDATE*
- *Windows IT Pro UPDATE*
- *Windows Tips & Tricks UPDATE*
- *WinInfo Daily UPDATE*

www.windowsitpro.com/email

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Diane Madzelonka at Diane.madzelonka@penton.com.

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the Web. Includes FREE access to eBooks and archived eLearning events, plus a subscription to either Windows IT Pro or SQL Server Magazine.
www.windowsitpro.com/go/vipsub

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

ASSOCIATED WEBSITES

DevProConnections

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePointPro Connections

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and Web seminars mentored by a community of peers and professionals.
www.sharepointproconnections.com

NEW WAYS TO REACH *WINDOWS IT PRO* EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: http://tinyurl.com/d5bquf. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Windows IT Pro

Xian
Wings 2010

Microsoft tech·ed Europe 2010 | Visit us at Stand E11

For BlackBerry®, Windows Mobile® and soon for iPhone®

System Center Operations Manager. Anytime. Anywhere.

Going Mobile? Take Microsoft® System Center Operations Manager with you.

Just another day out of the office. You're trying to watch a movie but your servers have different plans. One machine is complaining about drive space and another is having problems talking to the database. Both need to be fixed right away. Wouldn't it be nice to keep control of your IT systems anywhere and anytime you want?

With Xian Wings 2010 on your smartphone you can:

* Receive event notifications
* Review system and application status
* Build graphs for key performance data
* Run custom tasks
* Manage alerts

Get a free trial now!
here: www.jalasoft.com/wings

1219 Durham | Houston, Texas 77007 | Toll Free: 888 402-6717 |

Jalasoft

by Jason Bovberg

USER MOMENT OF THE MONTH

Years ago, I worked at a small agricultural-research firm in the Midwest. One day, I got a confusing call from a user, who asked, "Is something going on with the power? In the past few weeks, my computer has been randomly rebooting, right in the middle of my work!" He said the monitor would flicker, the whole system would shut down, and a reboot would happen almost immediately. I walked over to the user's office to see if I could observe the symptoms. At once, I heard the rather loud music coming from the iPod dock on his desk. I instructed him to just do his work as usual. He turned the music down a bit and worked for a while. After a few minutes, sure enough, the system powered off and rebooted. It was only after thoroughly checking cable connections that I realized the user was tapping his foot to the music coming from the iPod—and in the process jostling the power strip just enough to affect power to the computer.

—Alan

Tech Quote of the Month

In a *Financial Times* article, Nokia Executive Vice President Anssi Vanjoki compared mobile-phone manufacturers that switch to Android with little Finnish boys who "pee in their pants" to warm themselves in the winter. This rather vivid analogy illustrated Vanjoki's opinion that moving to Android is merely a short-term solution that will bring trouble later. Do you think he has a (soggy and cold) point?

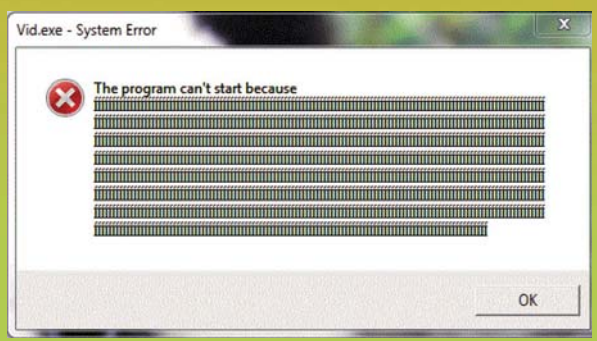

Figure 1: Obviously!

Figure 2: Wait, who's updating the Update updater?

Odd Product Announcement of the Month

In its announcement of the new VMware vSphere 4.1, VMware talks up the product's potency, reliability, and efficiency, but it does so with notably odd phrasing. Check out this announcement image, which promises that vSphere 4.1 "sparkles everything. Once again!" We're not entirely sure what that means, but it certainly sound zesty and fresh, with a nice lemony scent!

November 2010 issue no. 195, *Windows IT Pro* (ISSN 1552-3136) is published monthly. *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 221 E. 29th St., Loveland, CO 80538. Advertising rates furnished upon request. Periodicals Class postage paid at Loveland, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538. SUBSCRIBERS:** Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 221 E. 29th St., Loveland, CO 80538. Printed in the USA.

Kiss your antivirus bloatware goodbye

NEW Premium Version with Firewall

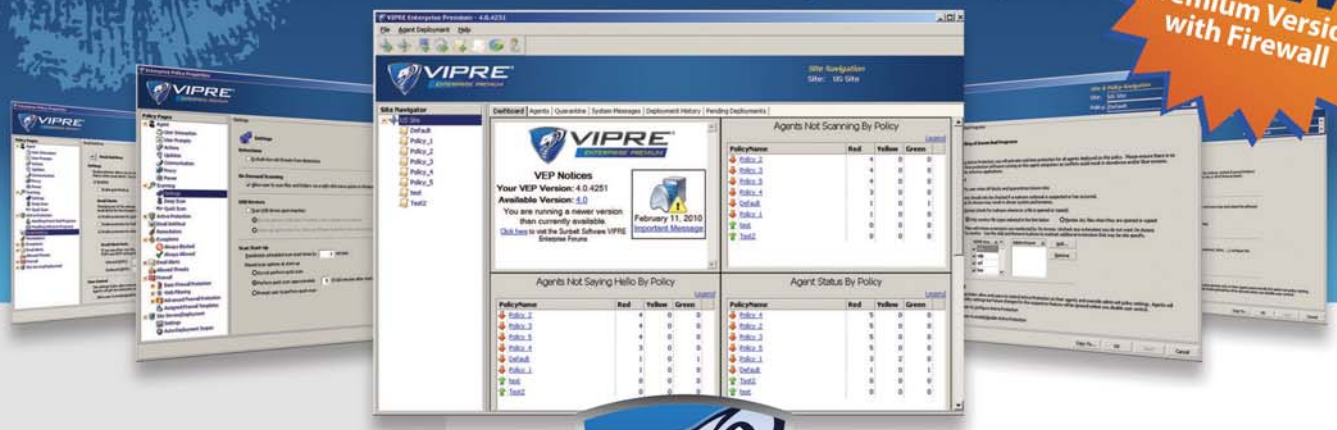

VIPRE®

TEST DRIVE ENTERPRISE PREMIUM

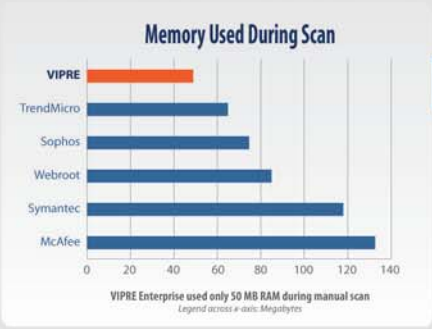

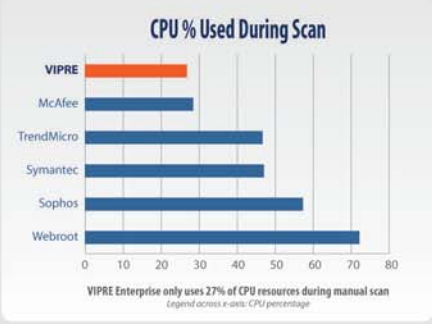

How does your current software compare? VIPRE Enterprise scans at a brisk 13.95 MB/sec and uses just 27% of CPU and 50 MB of RAM. In idle, it uses a mere 13.3 MB RAM with a disk footprint of just 113 MB. You'll hardly notice it's running!

Special Competitive Upgrade: *50% Discount!*

Until now, antivirus engines have been Frankensteins, bolted together from bits and pieces of different products. They're slow, full of bugs, and hard to manage.

VIPRE Enterprise Premium is a revolutionary new approach. It combines high-performance antivirus, antispyware, and desktop firewall into a single agent so you get comprehensive endpoint malware protection with low system resource usage. It's fast, powerful and easy.

Plus, advanced anti-malware technology protects your system against the new wave of malware threats. No more juggling multiple programs. No more dealing with user complaints about slow workstation performance.

- **COMPLETE!** All-in-one protection from today's malware.
- **FAST!** High-performance and low impact on system resources.
- **EASY!** Manage everything easily from one command screen.
- **RELIABLE!** Configurable, real-time monitoring technology.
- **AFFORDABLE!** Ask for a quote with our 50% competitive upgrade discount!

Why struggle with slow resource hogs when you can manage ALL your malware threats with one fast, easy application?

Curious? Download your FREE copy of VIPRE Enterprise Premium and give it a test drive.

When you compare VIPRE Enterprise Premium to Symantec, McAfee, Trend Micro or whatever antivirus program you're using, **you WILL want to switch!** Don't worry, though. You can get VIPRE Enterprise Premium with a **50% competitive upgrade discount!**

Compatible with Windows 7

BEST IN TECH 2009 Windows IT Administrator

Windows IT Pro 2009 Community Choice GOLD

vb 100 VIRUS virusbtn.com

Sunbelt Software
Part of the GFI Software Family

Plus we will buy out your existing maintenance contract for 1 year!

Download now: www.TestDriveVipre.com

Sunbelt Software Tel: 1-888-688-8457 or 1-727-562-0101 Fax: 1-727-562-5199 www.SunbeltSoftware.com sales@sunbeltsoftware.com

Being competitive starts with being productive. Here's your edge.

To help your business be more productive, the IBM® System x3650 M3 Express® server, featuring the Intel® Xeon® processor 5600 series, can help you achieve up to 40% increased performance[1] With more storage and memory capacity, it is now possible to access and process more data than ever before–helping you to efficiently meet your increased business demands.

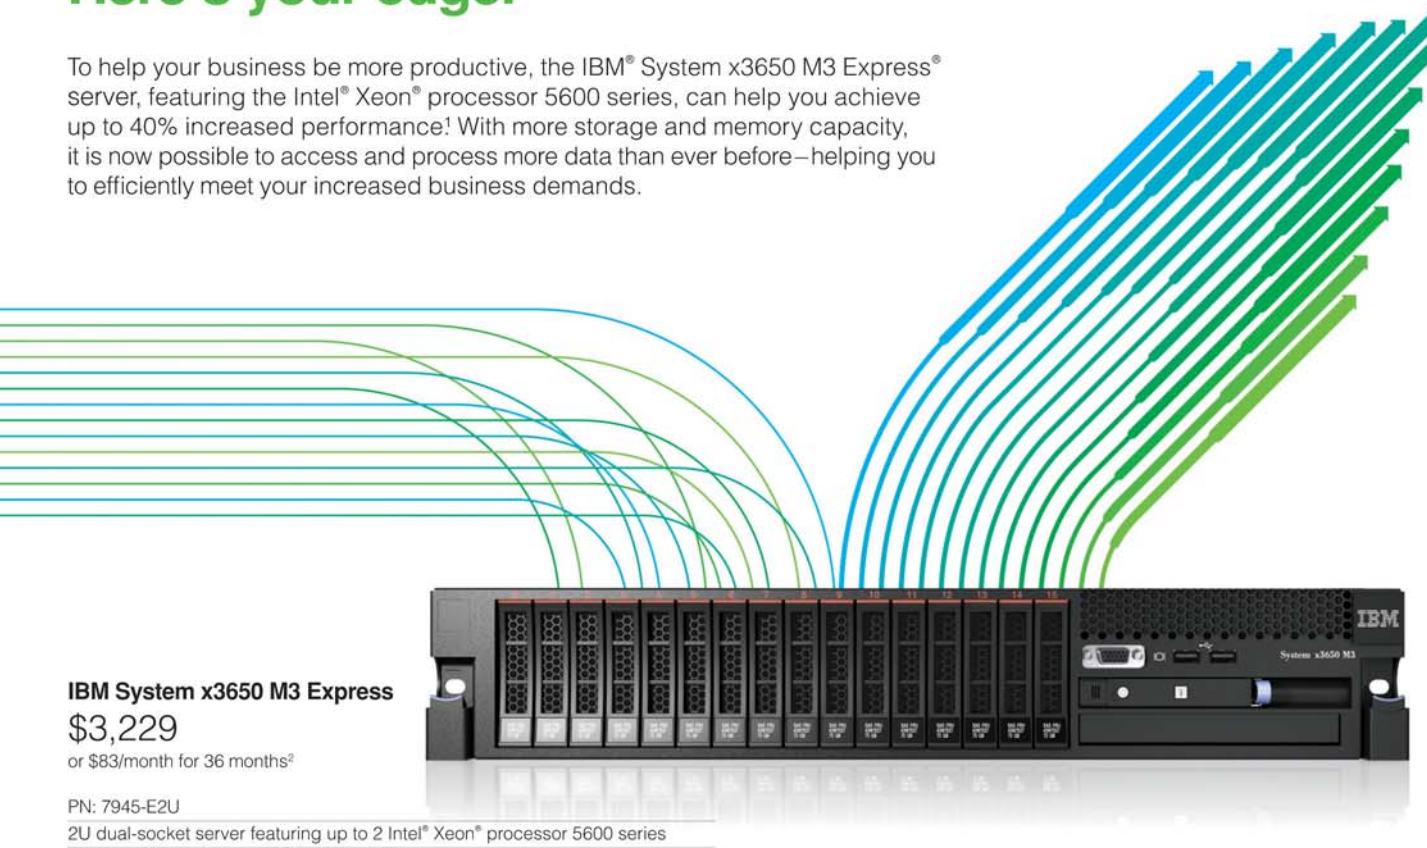

IBM System x3650 M3 Express

$3,229

or $83/month for 36 months[2]

PN: 7945-E2U

- 2U dual-socket server featuring up to 2 Intel® Xeon® processor 5600 series
- Energy-efficient design, 92% efficient PS
- 3 HS fan modules, altimeter

IBM System x3550 M3 Express

$1,969

or $51/month for 36 months[2]

PN: 7944-E1U

- 1U dual-socket server featuring up to 2 Intel® Xeon® processor 5600 series
- Energy-efficient design, 92% efficient PS
- 6 HS fan modules, altimeter

IBM System Storage® DS3500 Express

$8,799

or $226/month for 36 months[2]

PN: 1746-A2D or 1746-C2A

- External disk storage with 3 Gbps Serial Attached SCSI (SAS) interface technology
- Scalable up to 3.6 TB of storage capacity with 300 GB hot-swappable SAS HDDs or up to 9 TB of storage capacity with 750 GB hot-swappable SATA HDDs

See for yourself.

See how much you could be saving–in just minutes–with the IBM Systems Consolidation Evaluation Tool.

ibm.com/systems/productivity
1 866-872-3902
(mention 6N8AH30A)

[1]Based on Intel Engineering Study, January 2010 - performance increase comparing latest Intel Xeon processor 5600 series to previous generation - Intel Xeon processor 5500 series. See page 8, footnote 3 for more information: http://www.intel.com/Assets/PDF/prodbrief/323501.pdf. [2]Global Financing offerings are provided through IBM Credit LLC in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government customers. Monthly payments provided are for planning purposes only and may vary based on your credit and other factors. Lease offer provided is based on an FMV lease of 36 monthly payments. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice. IBM hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply. For a copy of applicable product warranties, visit http://www.ibm.com/servers/support/machine_warranties. IBM makes no representation or warranty regarding third-party products or services. IBM, the IBM logo, System Storage and System x are registered trademarks or trademarks of International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM trademarks, see www.ibm.com/legal/copytrade.shtml. Intel, the Intel logo, Xeon and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries. All other products may be trademarks or registered trademarks of their respective companies. All prices and savings estimates are subject to change without notice, may vary according to configuration, are based upon IBM's estimated retail selling prices as of 8/09/10 and may not include storage, hard drive, operating system or other features. Reseller prices and savings to end users may vary. Products are subject to availability. This document was developed for offerings in the United States. IBM may not offer the products, features, or services discussed in this document in other countries. Prices are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or IBM Business Partner for the most current pricing in your geographic area.